

Edited by
M. A. H. Dempster and Ke Tang

Commodities

Chapman & Hall/CRC FINANCIAL MATHEMATICS SERIES

VISIT...

LANZAROTE
Caliente.COM

VISIT...

LANZAROTE
Caliente.COM

Commodities

CHAPMAN & HALL/CRC

Financial Mathematics Series

Aims and scope:

The field of financial mathematics forms an ever-expanding slice of the financial sector. This series aims to capture new developments and summarize what is known over the whole spectrum of this field. It will include a broad range of textbooks, reference works and handbooks that are meant to appeal to both academics and practitioners. The inclusion of numerical code and concrete real-world examples is highly encouraged.

Series Editors

M.A.H. Dempster

Centre for Financial Research

Department of Pure

Mathematics and Statistics

University of Cambridge

Dilip B. Madan

Robert H. Smith School

of Business

University of Maryland

Rama Cont

Department of Mathematics

Imperial College

Published Titles

American-Style Derivatives; Valuation and Computation, *Jerome Detemple*

Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option

Pricing, *Pierre Henry-Labordère*

Commodities, *M. A. H. Dempster and Ke Tang*

Computational Methods in Finance, *Ali Hirs*

Counterparty Risk and Funding: A Tale of Two Puzzles, *Stéphane Crépey and*

Tomasz R. Bielecki, With an Introductory Dialogue by Damiano Brigo

Credit Risk: Models, Derivatives, and Management, *Niklas Wagner*

Engineering BGM, *Alan Brace*

Financial Mathematics: A Comprehensive Treatment, *Giuseppe Campolieti and*

Roman N. Makarov

Financial Modelling with Jump Processes, *Rama Cont and Peter Tankov*

Interest Rate Modeling: Theory and Practice, *Lixin Wu*

Introduction to Credit Risk Modeling, Second Edition, *Christian Bluhm,*

Ludger Overbeck, and Christoph Wagner

An Introduction to Exotic Option Pricing, *Peter Buchen*

Introduction to Risk Parity and Budgeting, *Thierry Roncalli*

Introduction to Stochastic Calculus Applied to Finance, Second Edition,

Damien Lamberton and Bernard Lapeyre

Monte Carlo Methods and Models in Finance and Insurance, *Ralf Korn, Elke Korn,*

and Gerald Kroisandt

Monte Carlo Simulation with Applications to Finance, *Hui Wang*

Nonlinear Option Pricing, *Julien Guyon and Pierre Henry-Labordère*

Numerical Methods for Finance, *John A. D. Appleby, David C. Edelman,*

and John J. H. Miller

Option Valuation: A First Course in Financial Mathematics, *Hugo D. Junghenn*
Portfolio Optimization and Performance Analysis, *Jean-Luc Prigent*
Quantitative Finance: An Object-Oriented Approach in C++, *Erik Schlögl*
Quantitative Fund Management, *M. A. H. Dempster, Georg Pflug,*
and Gautam Mitra
Risk Analysis in Finance and Insurance, Second Edition, *Alexander Melnikov*
Robust Libor Modelling and Pricing of Derivative Products, *John Schoenmakers*
Stochastic Finance: An Introduction with Market Examples, *Nicolas Privault*
Stochastic Finance: A Numeraire Approach, *Jan Vecer*
Stochastic Financial Models, *Douglas Kennedy*
Stochastic Processes with Applications to Finance, Second Edition,
Masaaki Kijima
Structured Credit Portfolio Analysis, Baskets & CDOs, *Christian Bluhm*
and Ludger Overbeck
Understanding Risk: The Theory and Practice of Financial Risk Management,
David Murphy
Unravelling the Credit Crunch, *David Murphy*

Proposals for the series should be submitted to one of the series editors above or directly to:
CRC Press, Taylor & Francis Group
3 Park Square, Milton Park
Abingdon, Oxfordshire OX14 4RN
UK

Chapman & Hall/CRC FINANCIAL MATHEMATICS SERIES

Commodities

Edited by

M. A. H. Dempster

University of Cambridge

and

Cambridge Systems Associates Limited, UK

Ke Tang

Tsinghua University, Beijing, China



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

A CHAPMAN & HALL BOOK

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2016 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20150522

International Standard Book Number-13: 978-1-4987-1233-0 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

Editors.....	xi
Contributors	xiii
Introduction	xvii

SECTION I Oil Products

1	Inconvenience Yield, or the Theory of Normal Contango	3
	<i>Ilia Bouchouev</i>	
2	Determinants of Oil Futures Prices and Convenience Yields	9
	<i>M. A. H. Dempster, Elena Medova, and Ke Tang</i>	
3	Pricing and Hedging of Long-Term Futures and Forward Contracts with a Three-Factor Model	31
	<i>Kenichiro Shiraya and Akihiko Takahashi</i>	
4	An Empirical Study of the Impact of Skewness and Kurtosis on Hedging Decisions	53
	<i>Jing-Yi Lai</i>	
5	Long-Term Spread Option Valuation and Hedging	71
	<i>M.A.H. Dempster, Elena Medova, and Ke Tang</i>	
6	Analysing the Dynamics of the Refining Margin: Implications for Valuation and Hedging	95
	<i>Andrés García Mirantes, Javier Población, and Gregorio Serna</i>	
7	Quantitative Spread Trading on Crude Oil and Refined Products Markets	119
	<i>Mark Cummins and Andrea Bucca</i>	

SECTION II Other Commodities

8	Inversion of Option Prices for Implied Risk-Neutral Probability Density Functions: General Theory and Its Applications to the Natural Gas Market	151
	<i>Yijun Du, Chen Wang, and Yibing Du</i>	

9	Investing in the Wine Market: A Country-Level Threshold Cointegration Approach	173
	<i>Lucia Baldi, Massimo Peri, and Daniela Vandone</i>	
10	Cross-Market Soybean Futures Price Discovery: Does the Dalian Commodity Exchange Affect the Chicago Board of Trade?	191
	<i>Liyan Han, Rong Liang, and Ke Tang</i>	
11	The Structure of Gold and Silver Spread Returns.....	213
	<i>Jonathan A. Batten, Cetin Ciner, Brian M. Lucey, and Peter G. Szilagyi</i>	
12	Gold and the U.S. Dollar: Tales from the Turmoil	227
	<i>Paolo Zagaglia and Massimiliano Marzo</i>	
13	A Flexible Model of Term-Structure Dynamics of Commodity Prices: A Comparative Analysis with a Two-Factor Gaussian Model	243
	<i>Hiroaki Suenaga</i>	
14	Application of a TGARCH-Wavelet Neural Network to Arbitrage Trading in the Metal Futures Market in China	279
	<i>Lei Cui, Ke Huang, and H.J. Cai</i>	

SECTION III Commodity Prices and Financial Markets

15	Short-Horizon Return Predictability and Oil Prices	301
	<i>Jaime Casassus and Freddy Higuera</i>	
16	Time-Frequency Analysis of Crude Oil and S&P 500 Futures Contracts	337
	<i>Joseph McCarthy and Alexei G. Orlov</i>	
17	Sectoral Stock Return Sensitivity to Oil Price Changes: A Double- Threshold FIGARCH Model	359
	<i>Elyas Elyasiani, Iqbal Mansur, and Babatunde Odusami</i>	
18	Short-Term and Long-Term Dependencies of the S&P 500 Index and Commodity Prices.....	387
	<i>Michael Graham, Jarno Kiviaho, and Jussi Nikkinen</i>	
19	Portfolio Selection with Commodities Under Conditional Copulas and Skew Preferences	403
	<i>Carlos González-Pedraz, Manuel Moreno, and Juan Ignacio Peña</i>	
20	Strategic Commodity Allocation	433
	<i>Pierre Six</i>	
21	Long-Short Versus Long-Only Commodity Funds	459
	<i>John M. Mulvey</i>	
22	Commodity Markets through the Business Cycle	469
	<i>Julien Chevallier, Mathieu Gatumel, and Florian Ielpo</i>	

23	The Dynamics of Commodity Prices	501
	<i>Chris Brooks and Marcel Prokopczuk</i>	
24	A Hybrid Commodity and Interest Rate Market Model	523
	<i>Kay F. Pilz and Erik Schlögl</i>	

SECTION IV Electricity Markets

25	Modelling the Distribution of Day-Ahead Electricity Returns: A Comparison	551
	<i>Alessandro Sapio</i>	
26	Stochastic Spot Price Multi-Period Model and Option Valuation for Electrical Markets	571
	<i>Eivind Helland, Timur Aka, and Eric Winnington</i>	
27	Modelling Spikes and Pricing Swing Options in Electricity Markets	585
	<i>Ben Hambly, Sam Howison, and Tino Kluge</i>	
28	Efficient Pricing of Swing Options in Lévy-Driven Models	607
	<i>Oleg Kudryavtsev and Antonino Zanette</i>	
29	Hedging Strategies for Energy Derivatives	623
	<i>Peter Leoni, Nele Vandaele, and Michèle Vanmaele</i>	
30	The Valuation of Clean Spread Options: Linking Electricity, Emissions and Fuels	645
	<i>René Carmona, Michael Coulon, and Daniel Schwarz</i>	
31	Is the EUA a New Asset Class?	667
	<i>Vicente Medina and Angel Pardo</i>	
Index		689

Editors

M.A.H. Dempster is professor emeritus at the Centre for Financial Research, Department of Pure Mathematics and Statistics, University of Cambridge. Educated at Toronto, Carnegie Mellon and Oxford Universities, he has taught and researched in leading universities on both sides of the Atlantic and is founding editor-in-chief of *Quantitative Finance* and the *Oxford Handbooks in Finance*. Consultant to many global financial institutions, corporations and governments, he is regularly involved in research presentation and executive education worldwide. He is the author of over 110 research articles in leading international journals and 14 books; his work has won several awards and he is an honorary fellow of the UK Institute of Actuaries, a foreign member of the Academia Lincei (Italian Academy) and managing director of Cambridge Systems Associates Limited, a financial analytics consultancy and software company.

Ke Tang is a professor in the Institute of Economics, School of Social Science, Tsinghua University, where he teaches courses in economics and finance. Before joining Tsinghua in 2014, he was a professor in the Hanqing Advanced Institute of Economics and Finance, Renmin University of China. He earned his BA in engineering from Tsinghua University in 2000, master of financial engineering from the University of California, Berkley, in 2004, and his doctoral degree in finance from Cambridge University in 2008. His research has covered such topics as commodity markets, Internet finance and Chinese stock markets. He has published many papers including *The Review of Financial Studies*, *Annual Review of Financial Economics*, *Journal of Banking and Finance*, etc. He is a frequent participant in various policy-related conferences held by institutions such as the United Nations Conference on Trade and Development, Organisation for Economic Co-operation and Development and the Food and Agriculture Organization of the United Nations. He currently serves as the managing editor of *Quantitative Finance*.

Contributors

Timur Aka

Axpo AG
Baden, Switzerland

Lucia Baldi

University of Milan
Milano, Italy

Jonathan A. Batten

Monash University
Caulfield East, Australia

Ilia Bouchouev

Koch Supply & Trading, L.P.
New York, New York

Chris Brooks

University of Reading
Reading, United Kingdom

Andrea Bucca

Pegaso Consultancy Services Limited
London, United Kingdom

H.J. Cai

Wuhan University
Wuhan, People's Republic of China

René Carmona

Princeton University
Princeton, New Jersey

Jaime Casassus

Pontificia Universidad Católica de Chile
Santiago, Chile

Julien Chevallier

IPAG Business School (IPAG Lab)
and
Université Paris (LED)
Paris, France

Cetin Ciner

University of North Carolina–Wilmington
Wilmington, North Carolina

Michael Coulon

University of Sussex
Brighton, United Kingdom

Lei Cui

Wuhan University
Wuhan, People's Republic of China

Mark Cummins

Dublin City University
Dublin, Ireland

M.A.H. Dempster

University of Cambridge, Centre for
Mathematical Sciences
and
Cambridge Systems Associates Limited
Cambridge, United Kingdom

Yibing Du

Nationwide Insurance
Columbus, Ohio

Yijun Du

Lower Colorado River Authority
Austin, Texas

Elyas Elyasiani

Temple University
Philadelphia, Pennsylvania

Mathieu Gatumel

University of Savoy – Institut de Recherche en
Gestion et Économie (IREGE)
Annecy, France
and

Centre de Recherches Appliquées à la Gestion
(CERAG)
Saint-Martin-d'Hères, France

Carlos González-Pedraz

Repsol Trading
Madrid, Spain

Michael Graham

Stockholm University
Stockholm, Sweden

Ben Hambly

University of Oxford
Oxford, United Kingdom

Liyan Han

Beihang University
Beijing, China

Eivind Helland

Axpo AG
Baden, Switzerland

Freddy Higuera

Universidad Catolica del Norte
Antofagasta, Chile

Sam Howison

University of Oxford
Oxford, United Kingdom

Ke Huang

Rutgers University
Newark, New Jersey

Florian Ielpo

IPAG Business School (IPAG Lab)
and
Université Paris I Panthéon-Sorbonne – Centre
d'Économie de la Sorbonne (CES)
Paris, France
and

Lombard Odier, Asset Management
Lausanne, Switzerland

Jarno Kiviaho

University of Vaasa
Vaasa, Finland

Tino Kluge

University of Oxford
Oxford, United Kingdom

Oleg Kudryavtsev

Russian Customs Academy Rostov Branch
and
Southern Federal University
Rostov-on-Don, Russia

Jing-Yi Lai

National Chung Cheng University
Taiwan, Republic of China

Peter Leoni

Katholieke Universiteit Leuven
Leuven, Belgium

Rong Liang

Renmin University of China
Beijing, China

Brian M. Lucey

Trinity College
Dublin, Ireland

Iqbal Mansur

Widener University
Chester, Pennsylvania

Massimiliano Marzo

Università di Bologna
Bologna, Italy

Joseph McCarthy

Bryant University
Smithfield, Rhode Island

Vicente Medina

Repsol Trading
Madrid, Spain

Elena Medova

University of Cambridge
and
Cambridge Systems Associates Limited
Cambridge, United Kingdom

Andrés García Mirantes

Instituto de Educacion Secundaria Candas
Candas, Spain

Manuel Moreno

Universidad de Castilla La Mancha
Toledo, Spain

John M. Mulvey

Princeton University
Princeton, New Jersey

Jussi Nikkinen

Stockholm University
Stockholm, Sweden

Babatunde Odusami

Widener University
Chester, Pennsylvania

Alexei G. Orlov

Radford University
Radford, Virginia

Angel Pardo

University of Valencia
Valencia, Spain

Juan Ignacio Peña

Universidad Carlos III de Madrid
Getafe, Spain

Massimo Peri

University of Milan
Milano, Italy

Kay F. Pilz

RIVACON GmbH
Friedrichsdorf, Germany

Javier Población

Banco de España
Madrid, Spain

Marcel Prokopczuk

Leibniz University Hannover
Hannover, Germany

Alessandro Sapio

Parthenope University of Naples
Naples, Italy

Erik Schlögl

University of Technology Sydney
Sydney, Australia

Daniel Schwarz

Carnegie Mellon University
Pittsburgh, Pennsylvania

Gregorio Serna

Universidad de Alcalá
Madrid, Spain

Kenichiro Shiraya

University of Tokyo
Tokyo, Japan

Pierre Six

Neoma Business School
Mont Saint Aignan, France

Hiroaki Suenaga

Curtin University
Perth, Australia

Peter G. Szilagyi

Central European University
Budapest, Hungary

and

University of Cambridge
Cambridge, United Kingdom

Akihiko Takahashi

University of Tokyo
Tokyo, Japan

Ke Tang

Renmin University of China
Beijing, People's Republic of China

Nele Vandaele

Ghent University
Ghent, Belgium

Daniela Vandone

University of Milan
Milano, Italy

Michèle Vanmaele

Ghent University
Ghent, Belgium

Chen Wang

Lower Colorado River Authority
Austin, Texas

Eric Winnington

Time-Steps AG
Affoltern am Albis, Switzerland

Paolo Zagaglia

Università di Bologna
Bologna, Italy

and

Rimini Centre for Economic Analysis
Rimini, Italy

Antonino Zanette

University of Udine
Udine, Italy

Introduction

Summary

In the past decade, global commodity markets have experienced dramatic growth in terms of both prices and trading volumes. Starting in the early 2000s, key commodity prices grew at rates unprecedented in history and then fell sharply during the financial crisis. However, from mid-2009, many commodity prices began climbing again, with oil prices increasing to over \$100 a barrel. The steady growth of U.S. oil production from hydraulic fracturing of shale, ‘fracking’, over this period and the slowdown of the growth of the Chinese economy have caused oil and iron ore prices to halve spectacularly at the time of writing—in spite of adverse geopolitical events affecting supply. Indeed, in both cases these remarkable events appear to have been caused by a mixture of supply and demand factors. A related phenomenon over the post-crisis period is that most commodity prices have become extremely volatile, actually much more volatile than would be implied by variations in the fundamental balance of demand and supply. The past decade has also seen a gradual increase in the financialization of commodities, which is likely related to this increased volatility of prices. By purchasing commodity-based exchange-traded funds (ETFs) or indices, an ever greater number of investors have begun to treat commodities as a new asset class for both portfolio diversification and inflation hedging. Because of high energy prices, some food crops have been used to produce biofuel over the period, so that fluctuations of energy prices have also influenced food prices. All these developments in commodity markets have of course coincided with the fast growth generally of emerging markets such as China, India and Brazil. The urbanization and expanding middle-class population of the emerging economies signals ultimately a strong global demand for commodities in the long-term future. In fact, the major income of more than 70 developing countries comes from exporting commodities. Therefore, price discovery and information transmission between commodity forwards and futures markets are key issues for the economic development of these countries. Furthermore, the linkages between different commodity markets and between commodity futures and other financial markets raise natural questions within the background of financialization and the use of biofuels. Closely related issues are the consequences of these developments for pollution control and global warming. All these topics are addressed in the chapters of this volume.

Introduction to Commodity Finance

Commodities have always been a part of our daily life and they are now generally also considered to have become an important asset class in financial markets. A *forward contract* between a buyer and a seller is an agreement for the seller to deliver a specified amount of a commodity to the buyer by a certain date in the future at an agreed price. The [figure](#) shows the nature of the first forward contracts used



in ancient Sumeria from 8000 BC.* Clay tokens used to represent the amounts of different commodities, including livestock, were baked into a hollow clay vessel for security after marking them in the soft clay on its outside before firing. Contract disputes at delivery time were resolved by breaking the vessel to find inside the representative tokens of the previously agreed amounts.

The term *forward* comes from contract law around 317 BC in ancient Athens. Modern forward contracts date back to those for rice in eighteenth century Japan and to those for wheat and other grains in the late nineteenth-century United States.

Most commodities, including gold, oil, natural gas, etc., are now traded by *futures contracts* in commodity futures exchanges which in current form date from the early twentieth century for those set up for agricultural products in Chicago, Illinois. These *derivatives* are *standardized* contracts, which involve a standardized amount of the underlying commodity to which they refer. (The contemporary importance of these first financial derivatives has recently been emphasized by the U.S. government with its placement of the development of regulation for *all* derivatives under the post-crisis Dodd–Frank Act in the hands of the Commodity Futures Trading Commission.) To prevent *default* at contract maturity by an out-of-the-money speculative counterparty—a common event in the nineteenth-century United States—prior to maturity the out-of-the-money party must periodically post *margin* with the exchange that is passed to the in-the-money counterparty when the current market price for the futures contract is in their favour. This is termed *marking to market*. The usual measure of current trading volume for a specific commodity on a futures exchange is termed *open interest*, which refers to the number of standardized contracts, usually including both futures and options on futures, for the underlying commodity that are currently outstanding, i.e., which have not been settled in the immediately previous time period.

An *option* is a financial derivative contract, which gives the holder the *right* but not the obligation to purchase (*call* option) or sell (*put* option) the underlying entity, here a commodity futures contract, at its maturity date (*European* option), or at any time up to its maturity (*American* option), at a fixed price termed the *strike* price. It is interesting to note that it was the empirical observation of agricultural

* We are indebted to Francesca Chiminello of Bloomberg for bringing our attention to the existence of these fascinating objects.

prices by Kendall that led to the random walk hypothesis for price changes (Kendall and Bradford Hill 1953), later made rigorous for futures contracts with stationary (explicitly not Brownian motion) prices by Samuelson (1965), leading to the efficient market hypothesis, geometric Brownian motion (GBM) and the theory of option pricing by Black, Scholes and Merton (Black and Scholes 1973, Merton 1973).

There are broadly two classes of traders in commodity futures markets: *hedgers* and *speculators*. Hedgers, who are normally producers, seek to hedge their future physical position in a commodity by selling futures contracts; speculators, pursuing speculative profits, take the buy side of these trades. Keynes (1923) argued that in order to attract speculators to take counterparty positions, hedgers need to offer *risk premia* to potential speculators by offering futures prices lower than they otherwise might be. These premia should increase with the maturity of the futures contract, leading to a *decreasing* futures *term structure*, i.e., the variation of futures price with contract maturity. This phenomenon is referred to as *backwardization* and the theory is termed that of *normal backwardation*. After Keynes, many researchers have attempted to test this theory, including Rockwell (1967), Chang (1985) and others. The results are mixed.

One important feature of commodities is *storability*, i.e., all commodities except electricity can be stored, even agriculturals, at least for some time. The *theory of storage*, advanced by Kaldor (1939), Working (1949), Brennan (1958) and others, proposes that *inventory* serves as a buffer to offset demand or supply shocks in commodity markets. When the market-wide inventory is at a relatively low level, which corresponds to a high demand for the physical commodity and a futures term structure in backwardation, holders of scarce inventory should derive a larger benefit from holding the commodity. The marginal benefit of holding one unit of the commodity relative to holding the corresponding futures contract is thus defined to be the *convenience yield* of the commodity, which corresponds to the dividend yield of holding a stock. The convenience yield represents the value of the option of using the stored physical commodity in production at any suitable time up to contract maturity. A larger convenience yield corresponds to a deeper backwardation of the futures term structure. Gorton et al. (2013) found that convenience yield is negatively proportional to the level of inventory of a commodity. Routledge et al. (2000) showed that the convenience yield relates to the probability of inventory stock-out. Convenience yield is, therefore, generally regarded as a distinguishing feature of commodities as an asset class. It is also the key to *modelling* the futures term structure. Gibson and Schwartz (1990), Schwartz (1997), Casassus and Collin-Dufresne (2005) and others propose a mean reverting Ornstein–Uhlenbeck process for the convenience yield in the framework of futures term structure modelling. The implication is that rather than the single factor model of Black (1976), until recently widely used by industry for trading commodity futures, at least a two-factor model, which includes a mean reverting convenience yield is needed to capture the term structure of futures contracts.

Since the early 2000s commodities have gradually been considered to be a new asset class. As a result, many institutions, including hedge funds, mutual funds and pension funds, began investing in commodities through commodity futures indices, ETFs, etc. This process of *financialization* is addressed in Tang and Xiong (2012), where it is argued that financial market factors influence commodity prices beyond the well-understood fundamental effects of supply and demand. For example, Tang and Xiong show that the return correlation between different commodities increased substantially from 2004 to 2008. Henderson et al. (2015) show that trades motivated for financial reasons lead to significant price movements of commodity futures.

At a more general level, the fundamental reference in the field of commodities is the monograph of Geman (2005), published a decade ago at about the beginning of the accelerated financialization of commodity markets noted in the previous paragraph. The 10 chapters of Section III of this volume represent the latest thinking in this development, in particular regarding the increased importance of oil and gas prices in equity markets.

A recent more theoretical monograph is that of Pirrong (2012), which describes a structural approach to storable commodity futures/forwards pricing in the spirit of Working, as opposed to the reduced form arbitrage-based approaches originated by Keynes and Kaldor in the 1930s. The *structural approach*

applies dynamic programming to derive the implications of non-negative inventories for prices, which are determined by a market trading equilibrium under rational expectations of the traders. The spectacular lack of empirical success with this approach using daily data is not surprising in that this theory essentially ignores the role of *speculation* and hence implies an *increasing* term structure of futures prices, which is termed *contango*.

Structure of the Volume

Having briefly set out the practical and theoretical background, we turn next to the contents of this volume. The chapters are mainly based on papers selected from two recent (2012–13) special issues of *Quantitative Finance* (QF) with the theme of commodity markets,* together with a few older and more recent papers from QF and one from the *Journal of Banking and Finance*. Topics include the fundamental theory of futures/forwards and derivatives pricing for major commodity markets, as well as the interaction between commodity prices and other financial markets and the real economy.

The book is divided into four sections of seven or more chapters each: *Oil Products*; *Other Commodities* (including natural gas, wine, soybeans, corn, gold, silver, copper, zinc and other base metals); *Commodity Prices and Financial Markets* (including stock, bond, futures and currency markets, index products and ETFs as well as the broad economy) and *Electricity Markets* (including emissions). Implicitly or explicitly, all the contributions to Sections I and II of this volume, which treat storable or agricultural commodities, take speculation into account through a consideration of markets over time being *either* in backwardation or contango. Up-to-date considerations of both trading and investment, discussed in the books of, respectively, Clark et al. (2001) and Till and Eagleeye (2007), are also considered in Sections I, II, and III. Section III, representing the recent financialization of commodity markets mentioned previously and consisting of 10 chapters, is appropriately the largest section of the volume. The chapters of Section IV treat current aspects of electricity markets and emission control previously considered in the two specialist monographs of Benth et al. (2008, 2014). All chapters contain the technical advances to be expected in contributions published in *Quantitative Finance*.

A brief outline of the contribution of each chapter in the four sections of the book follows.

Chapter Contributions

Section I Oil Products

Chapter 1, entitled Inconvenience Yield or the Theory of Normal Contango, by Bouchouev, argues that structural changes in the demand and supply for hedging services are increasingly determining the price of oil. As indicated above, there has traditionally been more producer than consumer hedging in commodity futures markets, not least for oil. While speculative investors used to cover this gap nicely, now the reasons for investment in oil futures have changed and demand for them far exceeds producers' requirements. Speculators no longer invest in the front end of the oil market primarily to collect a positive risk premium delivered by the roll yield of a futures term structure in backwardation. Instead, they are hedging against inflation, possible weakness of the U.S. dollar and the risk of geopolitical events that could decrease their investment portfolios, resulting in a futures term structure in contango. Recently it has been the oil *storage* operator who has been buying relatively cheap physical oil in the spot market from producers and selling futures to financial hedgers in order to collect the risk premium in the form of *negative* roll yield. This has led to the purchase of *calendar spread* options based on the spread

* We thank Collette Teasdale, formerly of Taylor & Francis, who pointed out the timely relevance of a collection of papers on commodity markets and made a substantial initial contribution to bringing this collection about.

between current and futures oil prices (calls for contango and puts for backwardization) becoming an increasingly popular risk management tool by investors. In short, these developments are the consequences of the financialization of the oil markets.

Chapters 2 and 3 treat latent (i.e., unobservable) three-factor affine Gaussian spot price models in state space form with the parameters of their assumed latent factor dynamics extracted from futures term structure data using the expectation–maximization (EM) algorithm and the Kalman filter.

Chapter 2, by Dempster, Medova and Tang, considers the question of what economic factors are drivers of the short-, medium- and long-term factors of these models? This is a question that has been treated by macroeconomists, but with financial markets treated as *residual* to the macroeconomy (see e.g., Ang and Piazzesi 2003). This approach is of course inappropriate for market modelling for pricing and hedging. Using weekly data for West Texas Intermediate (WTI) crude on the New York Mercantile Exchange (NYMEX) from mid-1986 to 2010, the authors estimate a structural vector autoregression model for the three latent factors using their three-latent factor estimates for WTI crude. They find that the short-term factor relates to demand variables in the physical commodity markets and to trading variables, such as the net short position of commercial hedgers, in the futures markets. Their medium-term factor relates to business cycle indicators, demand and trading variables. Since an affine combination of the mean reverting short- and medium-term factors is identified as convenience yield in their model, these findings accord with both the usual backwardization and inventory theories under different market and macroeconomic conditions. The long-term factor on the other hand relates mainly to financial market variables, as are studied in Section III of this volume.

Using a similar three-factor model, Shiraya and Takahashi study both oil futures and copper forwards markets, NYMEX and London Metal Exchange, respectively, in Chapter 3. The authors use weekly data from 1997 to 2007 inclusive for WTI crude and from September 2002 to the end of 2007 for copper. They compare the hedging performance of one-, two- and (their) three-factor models. In particular, they show how to hedge using three contracts and demonstrate the superiority of their hedge over a traditional single rolling contract hedge, using both time series and simulation methods. They also show that a three-factor model outperforms a two-factor model in replicating the actual term structure and that stochastic mean reversion models are more effective than fixed mean reversion models in out-of-sample hedging.

In Chapter 4, Lai investigates the improvement over the classical method of hedging a cash position in a commodity with a futures portfolio, i.e., choosing minimal hedge ratios using mean–variance optimization (see e.g., Wilford 2014), obtained by adding time-varying skewness and excess kurtosis of the underlying commodity returns in the minimization objective. The author uses a GARCH (Generalized Autoregressive Conditional Heteroskedastic) model to estimate the mean and variance of Brent crude returns, an ARCH (Autoregressive Conditional Heteroskedastic) model for their skewness and kurtosis and a modified normal distribution to estimate the parameters of the resulting hedge ratios using daily data from 2000 to 2008 inclusive for Brent crude oil futures on the Intercontinental Exchange. He finds that adding a preference for positively skewed returns to the objective may not lead to more speculative investment and hedging behaviours, as had previously been thought. Moreover, analysis of the hedged portfolio returns suggest that the third moments of hedged portfolios have probably been well adjusted by mean–variance strategies, which provide empirical support for these traditional strategies.

Dempster, Medova and Tang investigate the cointegration of commodity prices in long-run dynamic equilibrium in Chapter 5. They assert that the spread between the prices of two such commodities should be modelled directly and propose a two-factor model for the spot spread, which is used to develop pricing and hedging formulae for options on spot and futures spreads. These results are illustrated for the *location spread* between Brent blend and WTI crude oil, using NYMEX daily futures data from 1993 to 2004 inclusive, and the refining *crack spread* between heating oil and WTI crude, using NYMEX daily futures data from 1984 to 2004 inclusive. The authors conclude that the resulting spread option values implied are significantly different from those of standard models, but are consistent with practical observations.

Continuing this theme, Chapter 6, by Mirantes, Población and Serna, empirically analyzes the dynamics of refining margins for heating oil and gasoline and traces the implications for options

valuation and hedging of these products. Using weekly observations on refining margin from January 1997 to May 2009, the authors show that the prices of these products are not only cointegrated with crude prices but also they share common long-term dynamics, the implication being that refining margin risk derives only from short-term effects, which may be hedged with crack spread options based on modelling the spread directly.

In the last contribution to this first section of the volume, Cummins and Bucca focus in Chapter 7 on daily trading strategies exploiting these short-term effects for WTI and Brent crude oil and heating and gas oil-refined products. They consider a total of 861 spreads over the period 2003–2010 and develop a novel statistical arbitrage trading model with generalized stepwise procedures controlling data snooping bias. By aggregating upward and downward mean reversion they find profitable trading strategies with Sharpe ratios often exceeding 2, with average daily returns from several to over 50 basis points and holding lengths from 9 to 55 days. They see a collapse in the number of profitable strategies in the crisis and generally investigate trading strategy robustness to transactions costs.

Section II Other Commodities

Section II begins with an exercise by Du, Wang and Du in Chapter 8 to extract the risk neutral probability density function (PDF) of natural gas from the daily closing prices of the September 2007 option contract on the NYMEX. Based on the classical option pricing formula of Cox and Ross (1976), they impose smoothness conditions on the PDF using penalty techniques and reduce this originally ill-posed inverse problem to a least squares problem in non-negative variables, which provides an optimal trade-off between fit and PDF smoothness. The authors compare their approach with a variety of alternatives in the literature to find that it not only gives the best fit to the data but also gives a “significantly better” model resolution in that it can resolve the long-wavelength features of the densities away from the observations, as well as resolving the details of the estimated densities within the observed strikes.

The next two contributions treat markets related to agricultural commodities, wine and soybeans, respectively. In Chapter 9, Baldi, Peri and Vandone consider investment in the global consumer wine market involving the use of a threshold cointegration model and apparently the first trading use of the Mediabanca Global Wine Industry Share Price Index. Most previous wine market studies have been concerned with high-priced fine wines for connoisseurs. Several previous studies have demonstrated that the returns of regional wine portfolios are comparable to those of the relevant stock exchange, but the emphasis in this work is different. Namely, it attempts to use global equity indices to identify investment opportunities in wine price indices by examining their different speeds of adjustment to the long-run equilibrium. The goal is to allow traders to identify information inefficiencies, which can be exploited to construct profitable investment strategies by buying and selling wine producers’ stocks. The non-linear cointegration techniques used allow the two index prices to adjust differently to large and small deviations from the long-run equilibrium through a threshold representing a desired profit level. Market indices are examined for a variety of new and old wine producing countries, including Australia, Chile, China, France and the United States using daily data from January 2001 through February 2009 obtained from Datastream. The results confirm the existence of threshold cointegration between the wine and composite industry equity indices in all countries studied except Australia, where the cointegration relationship is linear. The authors also consider the implications of economic development, macroeconomic conditions, political system, and financial market sophistication for these results.

Chapter 10, by Han, Liang and Tang, analyzes the cross-market effects on global price discovery for soybean futures of the Dalian Commodity Exchange (DCE) in China and the Chicago Board of Trade (CBOT) exchange in the United States. Using daily data from March 2002 to September 2011 inclusive and structural vector autoregression and vector error correction econometric models for soybean futures in both trading and non-trading hours, the authors demonstrate for the first time that information transfers between the DCE and CBOT flow in *both* directions with similar price impact

magnitudes. Perhaps not surprisingly, this shows that the DCE plays a prominent role in global soybean futures price discovery, in spite of opposite findings in the earlier literature.

The remaining four chapters of this section of the volume concern precious and base metals. They mainly focus on gold and silver markets, but the last chapter discusses trading base metals on the Shanghai commodity exchange.

In Chapter 11, Batten, Ciner, Lucey and Szilagyi study the structure of gold and silver spread returns in terms of the dynamics of the bi-variate relationship between gold and silver prices. They analyze the spread between the daily prices of near-month futures contracts for New York Commodity Exchange (COMEX) gold and silver trading on the NYMEX from January 1999 to December 2005. After first filtering the spread return process using ARMA (Autoregressive Moving Average) techniques to remove the short-term memory effects observed in short-term autocorrelations, the resulting residual series reveals time-varying non-linearities, which are identified as fractality using Hurst techniques. Using rolling 22- and 55-day windows for the estimation of a local measure of the Hurst coefficient, long run time dependency of this series after volatility rescaling is found to be significant by the classical Hurst–Mandelbrot–Wallis tests. The authors go on to test the effectiveness of simple trading rules based on the Hurst coefficient and find these superior to both the simple moving average strategies commonly applied by traders and buy and hold.

Chapter 12 turns to the impact of U.S. dollar (USD) currency fluctuations on the gold price through the crisis. In this chapter, Zagaglia and Marzo study the dynamic relationships between the spot price of gold and the USD–EUR and USD–GBP exchange rates with the euro and the pound, respectively, using daily data over the period from 13 October 2004 to 5 March 2010. They find that the onset of the turmoil in August 2007 had little effect on these variables until the Lehman bankruptcy in September 2008, and that the gold price has been more stable through the crisis than these USD exchange rates. They conclude that gold prices are less affected by market uncertainties than the USD, implying that gold can be a valuable addition to currency portfolios. Moreover, they suggest that holding gold, and selected other currencies, in a portfolio removes the risks of the transmission of volatility shocks caused by fluctuations in USD rates, independent of the severity of market fluctuations.

Suenaga compares two approaches to modelling the term structure of commodity futures prices in Chapter 13. The author specifies a two-factor (log) futures contract price model in terms of two time varying conditionally heteroskedastic latent factors and an idiosyncratic term, which are based on three deterministic functions of the contract duration and maturity date. This allows the volatility of daily futures returns to be captured directly. The resulting specification is sufficiently flexible to capture seasonal variation and complex contemporaneous volatility effects across contracts and maturities. These are usually missed by the standard two-factor Gaussian spot price model to which it is compared (see also Tang (2012) in this context). Since both models contain two latent factors, EM algorithm methods alternating Kalman filtering and maximum likelihood must be used to estimate their parameters. The comparison of the two alternative models and their composite, using daily futures data for crude oil, natural gas, corn and gold from the late 1980s to end 2007, shows that the new flexible model is superior for hedging and in fitting the complex realized dynamics of futures prices. In particular, the flexible model results show that price volatility exhibits a strong maturity effect for crude oil, natural gas and corn and strong seasonality of demand for natural gas and supply of corn. The conventional two-factor Gaussian model in state space form under standard uncorrelated serially independent measurement errors cannot model these complex dynamics (in this regard see also Dempster and Tang (2011) who consider copper prices).

The last chapter of Section II, like the last chapter of Section I, considers daily statistical arbitrage trading, here of zinc contracts on the Shanghai Futures Exchange. However Cui, Huang and Cai use entirely different methods in Chapter 14 to develop trading strategies based on a threshold GARCH-wavelet neural network. Four months of 2010 tick data are used to establish the optimal training and trading thresholds and the new method is compared to strategies based on an historical projection technique and a standard backpropagation neural network. Robustness checks on the rates of return of

the three methods are performed under different market conditions to show the clear superiority of the new method. Finally, the authors investigate the effects of trading commission charges on the returns of their method and show that for current commission levels their method remains very profitable. They also note that their profitable trading strategy is risky in that it has difficulty coping with sudden large price changes and should therefore be used in conjunction with a suitable risk management system (see e.g., Dempster and Leemans 2006).

Section III Commodity Prices and Financial Markets

The 10 chapters of this third section of the volume investigate contemporary aspects of the financialization of commodities. The first three chapters, 15, 16 and 17, consider the impact of oil prices on U.S. equities, while Chapter 18 widens the consideration of this impact to a range of commodity index drivers in a search for commodities whose prices are uncorrelated with equity prices and so can diversify investors' portfolios. This theme of investment portfolio commodity enhancement is further investigated from different viewpoints in Chapters 19, 20, 21 and 22. The final two chapters of this section, Chapters 23 and 24, investigate shorter-term models for derivative pricing.

In Chapter 15, Casassus and Higuera show that oil price shocks in terms of three-month futures returns are a strong predictor of U.S. excess stock returns at short horizons. They use as their shock variable quarterly WTI oil futures returns on the NYMEX from Q2 1983 (when trading in these contracts began) through Q4 2009 and the corresponding U.S. stock returns were obtained from the value-weighted index of the Chicago University Center for the Study of Security Prices (CRSP) over the same period. The authors demonstrate that oil price shocks are a non-persistent negative leading indicator of excess equity returns whose impact dies away after three quarters. They compare their futures-based oil price shock variable with other popular predictors based on quarterly data over a sample period, which includes the crisis, to suggest that oil price change is the *only* variable with (linear) forecasting power for stock returns which, as they demonstrate, is robust against the inclusion of other macroeconomic and financial variables in out-of-sample tests. The authors also study the cross section of expected stock returns using an oil price shock conditional CAPM (Capital Asset Pricing Model), which is highly statistically significant and a better fit to the data than the alternative conditional and unconditional models considered, including the Fama–French three-factor model. The leading indicator for excess stock returns they propose has the practical advantage over others of being market based, observable (rather than involving an arbitrary and possibly changing technical construction) and readily available to investors.

McCarthy and Orlov perform in Chapter 16 a time–frequency analysis of the relationship between daily closing prices of near month futures contracts on WTI crude and the S&P 500 using data from end March 1983 to beginning March 2011. The authors use frequency domain techniques—wavelets and cross-spectra—to examine the association between crude oil and equity prices in order to reveal insights into this relationship that are not apparent from a conventional time domain approach. They find a *long-term positive* correspondence between contemporaneous oil and S&P 500 prices in levels. This of course does not contradict the findings in the previous chapter of a *short-term negative* lagged relationship between oil price returns (shocks) and excess returns of the S&P 500. Their results also confirm that, when analyzed separately, oil prices lead oil volume and S&P 500 trading volume leads S&P 500 prices. These findings persist over a large number of time scales and across a wide range of Fourier frequencies. Their use of frequency domain techniques overcomes the simultaneity problems which have plagued some earlier investigations of the impact of oil prices on the U.S. economy. Taken together, the practical implications for investors of this chapter and the previous one are profound.

The next chapter, Chapter 17, by Elyasiani, Mansur and Odusami, investigates the relationship between oil price changes and U.S. stocks at the sectoral level with particular regard to volatility persistence, which affects asset pricing to reinforce and enhance the cross-sectional findings in Chapter 15. A non-linear double-threshold Fractionally Integrated Generalized Autoregressive Conditional

Heteroskedastic (FIGARCH) specification of sectoral relationships is used because it nests the GARCH and IGARCH (Integrated Generalized Autoregressive Conditional Heteroskedastic) models to allow the identification of the true extent of long memory in conditional return volatility at the sector level. The authors use daily one-month futures returns from January 1991 to December 2006 as oil price shocks and daily sector equity index data from the CRSP and Kenneth French databases on 10 U.S. industry sectors broadly classified as oil substitutes (2), oil-related (2), oil consuming (5) and financial (1). They consider different conditional oil return volatility regimes, calm and volatile, to adjust the FIGARCH thresholds for a Fama–French model augmented by oil returns when oil return shocks become large enough to exceed existing thresholds. The symmetry and signs of these thresholds are seen to vary widely by industry, with those of the oil-related industries found to be the tightest, i.e., most sensitive to volatility. The decay patterns of industry responses to oil price shocks are also seen to vary widely, which call for different policy responses by oil producers, oil consumers and regulators to these shocks. These sector-specific and period-specific results are tested against simpler model alternatives, which are found wanting. The chapter concludes that regime shifts and long-term structural changes due to oil price shocks have profound influence on how investors perceive long-term risk and form their expectations.

Chapter 18, by Graham, Kiviahio and Nikkinen, investigates the difference between short-term and long-term dependencies between a classification of all global commodity and S&P 500 returns. This difference is important to the dissimilar goals of short- and long-term equity investors. The authors use a wavelet coherency methodology to simultaneously take into account time and frequency domains in a manner similar to Chapter 16 for oil price shocks. They address the question of which, if any, commodity returns can genuinely be considered uncorrelated with equity returns and thus be useful to investors in actually diversifying their portfolios. The study uses weekly observations from January 1999 to December 2009 inclusive on the S&P 500 and the S&P GSCI (Goldman Sachs Commodity Index) and its 10 subcategories: Energy, Light Energy, Non-Energy, Reduced Energy, Agriculture, Livestock, Petroleum, Industrial Metals, Precious Metal and Softs. This sample period allows an analysis of the effects of the onset of the financial crisis on the studied relationships and a distinction between short-term (high-frequency) and long-term (low-frequency) effects. The findings show negligible short-term effects of the crisis on commodity–equity relationships but stronger, although still quite weak, long-term effects, as is commonly asserted. The authors conclude that some important distinctions can be made for combining equity and the various segmented commodity indices in investors' portfolios. In agreement with Chapter 11, they note in particular that the diversification benefits of combining equity (S&P 500) returns with precious metal (S&P GSCI Precious Metal) returns are robust for both short- and long-term investors across time.

The next four chapters of this volume consider various aspects of portfolio construction involving commodity components. In Chapter 19, González-Pedraz, Moreno and Peña investigate the portfolio selection problem of an investor in commodity futures with three-moment preferences in terms of mean, variance and skew of return distributions. This study augments that of Chapter 4, for daily oil futures of multiple contract maturities, using weekly near month futures data for oil, gold and equities from June 1990 to September 2010, respectively, for WTI crude from the NYMEX, for gold bars from the COMEX and for the S&P 500 equity index. The authors use a multi-variate dynamic t conditional copula with GARCH evolution of its dependence matrix and marginal conditional returns using a model with modified Student t innovations and autoregressively evolving parameters. They show that their statistical model fits the in-sample period data from 1990 to 2006 better than standard alternatives. For portfolio construction using this statistical model, they maximize a three-term portfolio objective, adding skew to positive return and negative return variance, with a non-linear budget constraint allowing short selling. Out-of-sample performance using one-step ahead forecasts of the marginal and copula parameters is measured in terms of Sharpe, Sortino and Omega ratios and exceeds that of the equally weighted and standard mean–variance optimal portfolios for their statistical model and variants. The study concludes that the specification of higher moments in the marginal return distributions and the type of their tail dependence has significant positive implications for out-of-sample portfolio allocation

performance with a weekly time step. This does not directly contradict the opposite conclusion reached in Chapter 4 for daily data regarding oil futures of various maturities, but it does raise questions for further work.

A first step in this direction is made in Chapter 20 by Six, who extends research into the financialization of commodities by considering investments in the term structure of some commodities individually, oil, copper and gold, as possible enhancements of a dynamic mean–variance efficient U.S. Treasury bond and equity strategic level portfolio. The model used to address this question is that of earlier work by Brennan, Schwartz and Lognardo (1997), who considered enhancement of such a standard strategic portfolio with U.S. Treasury bond futures. Using weekly data from September 2009 to March 2014, the author estimates the parameters of relatively standard diffusion models for stock price index (S&P 500), short rate (U.S. Treasury bonds) and commodity spot price and convenience yield (Chicago Mercantile Exchange futures), but with affine market prices of risk for equity and the commodity. The strategic investor has a constant relative risk aversion utility function, which is applied to maximize terminal wealth at a finite horizon. Without commodity investment, a dynamic mean–variance investment is optimal and is subsequently used as a benchmark over the data period. Portfolios to which are added the spot commodity and the spot commodity plus a long-dated commodity futures contract are compared to the benchmark. This in-sample empirical study over the crisis period reveals that the optimal commodity investment strategy for oil is a calendar spread, while that for copper and gold is a reverse calendar spread with offsetting positions in two highly correlated returns. For gold, these offsetting positions are extreme and nearly equal. As a result, when measured by the Sharpe ratio or terminal wealth, investment in any one of the spot commodities improves in-sample portfolio performance over the benchmark, but additional investment in their long-dated future does not, and may even seriously degrade portfolio performance. It is perhaps a pity that no out-of-sample backtest is available to augment these results, but one could expect in this setting that out of sample the inefficiency of long-dated futures is even greater based on the current post-crisis sample period. The author notes that his findings extend those of Brennan and Schwartz for an earlier period to reach an *opposite* conclusion. This opposite conclusion does not, however, contradict that of Chapter 4. Indeed, even though highly dynamically risky asset return volatilities (resulting from the dynamic volatilities of the stochastic market prices of risk employed in this study) lead to highly volatile weights for the commodity in the much higher yielding dynamic strategic portfolios, the weights for the stock and bond indices remain close to the benchmark weights for the mean–variance frontier.

In a much shorter contribution, Chapter 21, Mulvey continues the theme and findings of the previous chapter to outline the construction of an investible index long–short commodity fund, which is meant to be used as a risk reducing overlay to a standard portfolio robust to all market conditions. After discussing the sources of alpha in commodity portfolios and introducing the *ulcer index* as the cumulative sum of the standard deviation of drawdown values, he sketches the outline of a commodity overlay portfolio designed to minimize the effects of drawdown risks and negative rollover risks when markets are in contango. Six established tactics are used, which carefully balance long and short positions in the overlay portfolio: long and short momentum, long and short futures curve, trend following and breakout. Each of these strategies is equally weighted in the portfolio and the long and short positions are balanced to be market neutral. However, noting that three market regimes, growth, transition and contraction, occur in historical respective proportions two-thirds and one-sixth each over time, the portfolio performance can be enhanced by regime detection and a tilt of about 20% in the appropriate long or short direction. Regular periodic rebalancing is also mentioned as a source of portfolio return enhancement (see e.g., Dempster et al. 2011). Unfortunately, neither the commodities nor the backtest results used in the implementation of these ideas by a major global exchange are presented in the chapter.

Chapter 22, by Chevallier, Gatumel and Ielpo, points out that if commodity markets can be shown to be strongly related to the regimes of the business cycle noted in the previous chapter as important for a long–short commodity overlay portfolio, then this evidence contradicts the widespread intuition that commodity markets are a strong source of diversification for a standard long only cash–bond–equity portfolio. Based on a Bloomberg daily data set from 1990 to 2012 for the United States,

European Monetary Union (EMU) and China, the authors comprehensively investigate this question by first evaluating the impact of news on an extensive set of commodity markets and then using a Markov regime-switching model to analyze the effects of U.S. business cycle regimes on global commodity markets as an asset class. They analyze news impacts using EGARCH (Exponential Generalized Autoregressive Conditional Heteroskedastic) regression techniques on markets for gold, silver, platinum, aluminium, copper, nickel, zinc, lead, WTI crude, Brent crude, gas oil, heating oil, natural gas, corn, wheat, coffee, sugar, cocoa, cotton, soybeans and rice to reveal that the response of commodity prices to economic surprises is strong during global downturns, but weak during expansion periods. They note that their results suggest that markets overreacted to news during the 2008–2009 crisis, as they did in 2001, but showed a decreased sensitivity to news flow in H2 2009–2010. Unsurprisingly, they note an increased sensitivity to economic activity in China over the whole period. The chapter goes on to estimate a Markov business cycle regime-switching model with five regimes to the economies of the U.S., EMU and China. From an investor's viewpoint, the authors' findings may be summarized as follows. In agreement with Chapters 18 and 19, investors should not add commodities to portfolios on the grounds of low correlation of a commodity's returns with those of a standard portfolio, unless they are at least prepared for the long term, since there is a strong correlation of commodity markets with risky assets during economic downturns. Secondly, economic influences on commodity markets are complex and vary greatly for the United States, Europe and China. Thirdly, a cyclical rotation amongst commodity markets over economic regimes is suggested; during strong growth, industrial metals and energy should outweigh agricultural and precious metals in portfolios and the reverse should apply in downturns. Finally, relevant to the findings of Chapter 20, the authors note that in 'stalled' periods of the U.S. economy, commodities generally outperform the S&P 500 with the obvious implications for portfolio construction in this regime of the cycle.

The final two chapters of this section of the volume investigate models for option pricing, involving, respectively, jumps and stochastic volatility for commodity options in Chapter 23, and the relationship between WTI crude prices and interest rates for commodity and interest rate options in Chapter 24. In Chapter 23, Brooks and Prokopczuk use a basic non-negative mean-reverting square root process to model log spot prices to which they add stochastic volatility à la Heston and jumps in both drift and stochastic volatility. They estimate the parameters of this succession of models, including three separate stochastic volatility models, using Markov Chain Monte Carlo (MCMC), which is a Bayesian technique used here with non-informative priors. MCMC breaks down the high-dimensional posterior parameter distribution of a model into its low-dimensional complete conditional parameter distributions and latent factors, e.g., volatilities, jump times and jump sizes, which may be sampled efficiently. The authors use U.S. daily spot price data for crude oil, gasoline, gold, silver, soybeans and wheat from 1985 to March 2010, obtained from the Commodity Research Bureau, and the S&P 500 from Bloomberg for comparison. Individual commodity analysis identifies jumps due to, for example, Organization of the Petroleum Exporting Countries increasing oil supply in 1986, the Gulf War in 1990, Hurricane Katrina in 2005 and the Lehman default in 2008. As expected, return correlations are found to be highest between commodities in the same market sector and, with some interesting exceptions, so are volatility and jump correlations, but there is little interaction found between commodity market sectors. The results demonstrate that return correlations for their models between commodities and the stock index are also low, but, more importantly for portfolios of securities with embedded contingent claims that their volatility correlations are low as well so that commodities may act as a volatility diversifier for such embedded derivative portfolios. Some evidence is found for simultaneous equity and oil or gold jumps, but these results vary with the model used. The alternative spot price models are compared for each of the six commodities by pricing European calls and barrier options and by the effectiveness of their delta hedges. Substantial differences are found across commodities and between models. The authors conclude that this indicates the heterogeneous nature of the broad commodity asset class, agreeing with the findings of Chapters 18, 20 and 22, and they stress the importance of choosing the appropriate model for each individual commodity and purpose at hand.

Chapter 24, by Pilz and Shlögl, specifies a joint model of commodity price and interest rate risks analogous to the fixed income multi-currency LIBOR market model by directly addressing the issues arising in applying this model to commodities. WTI crude oil is used as the ‘foreign currency’ and USD as the domestic, with the spot ‘exchange rate’ as WTI near month futures price and the WTI futures term structure used for ‘foreign’ forward rates. The difference between futures and the forward rates used in fixed income must be specifically taken into account, by converting the futures term structure to an appropriate forward rate term structure by a suitable convexity correction. The authors construct a procedure to calibrate the model to market data for interest rates and commodity futures prices. NYMEX futures prices for 5 May 2008 are used to give the underlying WTI ‘forward’ data. Calibration of the model is to options on these futures prices and to caplets, caps and swaptions on U.S. forward rates, together with three-month prior historically estimated correlations between interest rates and futures prices. Standard Kirk approximations are used to *simultaneously* estimate prices for calendar spread options on futures prices and rates, which are favourably compared in-sample with Monte Carlo simulations from the calibrated model. The authors conclude that the main contribution of their model is in its calibration, which they feel is appropriate to other settings.

Section IV Electricity Markets

The seven chapters of this final section of the volume give an overview of the current and future modelling of electricity markets. Chapters 25 and 26 model electricity prices and demonstrate their complexity, while Chapters 27 and 28 treat swing options, the major derivative product used to mitigate electricity price risks. Chapter 29 considers hedging electricity cash flow risks more generally. Finally, Chapter 30 discusses clean spread options, which link electricity, fuels and emission allowances, while, in the spirit of Section III, Chapter 31 considers European Union Allowances (EAUs) for carbon emissions as a possible asset class.

In Chapter 25, Sapio investigates the statistical modelling of day-ahead returns in the NordPool (Scandinavia), APX (Netherlands) and Powernext (France) European electricity markets. He points out that in these markets 24 auctions are run simultaneously each day to set prices for each hour of the following day. The day-ahead prices are determined by uniform price auctions so that *all* power is purchased and sold for the next day at the market-clearing prices. Using daily data from these three principal markets, NordPool from 1997 to 2002 inclusive, APX from 2001 to 2004 and Powernext from 2002 to 2007, the author compares the goodness-of-fit to empirical densities over these periods for four alternative return densities, alpha-stable, normal inverse Gaussian (NIG), exponential power and asymmetric exponential power, using Kolmogorov–Smirnov and Cramer–von Mises statistics. The author finds that electricity returns display heavy tails across both markets and hourly auctions; however, returns are defined (log, percentage or price change) and seasonality is accounted for. The tails in the APX market and in daytime auctions are fatter than in other markets and in night-time auctions, but are dampened by the logarithmic transformation (calling into question its widespread use in financial economics of all sorts). Skew is an essential of *all* the fitted distributions, but the alpha-stable and NIG distributions systematically outperform the other two in goodness-of-fit to the data, although no clear ranking between these two emerges. Since both of them are closed under convolutions, daily return fat tails will persist for returns over longer intervals. The fitted alpha-stable distributions have exponents between 1 and 2, so that only the empirical *average* returns exist, but the empirical *variance* does not converge, unless of course the NIG distribution provides the better fit to a specific data set. Some potentially interesting cross-market comparisons are inconclusive, using percentage or price change returns, but log returns indicate that skew is negative in Powernext, but positive in APX. The study reveals some interesting intraday patterns that are not observable when daily returns are averaged. The author sees these results as a starting point for further work, particularly as regards policy for retail market liberalization and the carbon emissions allowance trading discussed in the last two chapters of this volume.

Helland, Aka and Winnington develop a multi-period model for electricity market spot prices and apply it to option pricing in Chapter 26. At the largest energy company in Switzerland, the authors have developed the model described in this chapter for the purposes of optimizing power plant production from pumped storage hydro parks, pricing complex options and structured products, and dynamically hedging production. They assert that a deep understanding of stochastic price dynamics is necessary for both risk management and derivative pricing, so that their model is built up in several steps and calibrated using spot price distributions, log returns of forward contracts and standard option contracts over a number of delivery periods. The base price level of their stochastic spot price model is a forward curve for each hour of the day, which has been estimated using a regression model and European Exchange (EEX) data for Germany. The main explanatory variables of this model are calendar dates (weekdays, holidays, etc.) to obtain (intraday, weekly and yearly) seasonality effects, historical meteorological data (daily temperature, wind and precipitation) and long-term market expectations. Each hourly price forward curve is then scaled to equal the power futures contract (PFC) price at each maturity to obtain an arbitrage-free model prediction. Daily changes in each PFC are based on a log-normal mean-reverting (Ornstein–Uhlenbeck) weekly process for the front-year forward base contract to reflect shifts in spot prices due to changes in long-term contracts. Similarly, a log-normal mean-reverting hourly process is used to model intraday and day-ahead perturbations in spot prices. Further, two jump processes are added in order to model the positive and negative price jumps observed in the German electricity market. Based on the most recent PFCs and a calibrated set of weekly, hourly jump processes, a set of price scenarios is generated by Monte Carlo simulation using a dynamic spot price model and then used to evaluate the price of options on futures available on the EEX. The fair prices of calls and puts for selected strikes are calculated from the scenarios and compared with observed prices for options on the EEX or with comparable broker data. The stochastic spot price multi-period model is recalibrated so that the monthly, quarterly and annual contracts fall within their quoted bid–ask spreads. The result is a practical and flexible model, which is actually used for all the purposes intended in an arbitrage-free manner.

In Chapter 27, Hambly, Howison and Kluge address the modelling of the price spikes occasionally observed in electricity market prices and use the result to price the swing options used by energy traders to manage the risks of these high prices. Swing contracts are a broad class of path-dependent options allowing the holder to exercise a certain right multiple times over a specified period, but exercising only one right at a time, with a ‘refraction period’ lag between exercises, or per time interval, e.g., once a day. For example, the right could be to receive a call or put option payoff, or a mixture of these, or one or the other with different strikes. Another popular feature allows the swing option holder to exercise a number, termed the volume, of calls or puts at once, typically with bounds on the volume and the sum of all multiple exercises. Such contracts protect the purchaser against excessive rises in electricity prices and, bundled with forward contracts, ensure a constant stream of energy at a predetermined price. If the strike price of the call options in a swing contract is set at a specific maturity forward price, the swing contract allows flexibility in the volume the holder will receive at the fixed price at that date. Such contracts can either ‘swing up’ or ‘swing down’ the volume received at forward dates, and hence the name. To price these contracts the authors use a general payoff function and an electricity spot price model that is the exponential of the sum of a deterministic seasonality function, a mean-reverting Ornstein–Uhlenbeck process and an independent mean-reverting pure jump process with a constant intensity and independent identically distributed jumps. The mean reversion speeds of the two stochastic components may be different, a useful feature for modelling NordPool contracts. Although this Scandinavian exchange and the UKPX (U.K. Power Exchange) and German EEX exchanges are mentioned, no explicit calibrations to forward curve data are given in the chapter. The authors derive the moment-generating function, and various approximations to the density, of the logarithm of the spot price process at a maturity date. They first present semi-analytic formulae for path-independent options and approximations for call and put options on forward contracts, with and without a delivery period. Then they use a tree/grid method of Jaillet et al. (2004) and an approximation to the conditional density of the spot price

process to price swing option payoff structures in the presence of price spikes obtained from the pure jump component of their spot price model, for which some numerical results are given.

The theme of pricing electricity swing options with multiple exercise rights is continued in Chapter 28 by Kudryavtsev and Zanette, using different models and methods. They model the spiking spot price process with a general Lévy process, which punctuates infinitesimally frequent small jumps with occasional large spikes. The authors concentrate on methods and numerical procedures which they illustrate with a swing option involving a fixed number of exercises of a futures put option within a specified maturity and with a positive refraction time. This multiple stopping time problem can be reduced to a cascade of single stopping time problems using ideas of Carmona and Touzi (2008), who used Monte Carlo methods for their numerical solution. The authors evaluate this problem using two different deterministic numerical techniques. They propose a Wiener–Hopf factorization method, which solves the system of integro-partial differential equations corresponding to the cascade and compare their approach to a more standard finite difference technique to solve the same system. Numerical results are given for pricing the swing option for both GBM and the Carr–Geman–Madan–Yor Lévy spot price processes, which show the reliability and robustness of both solution methods. Although again no empirical results are given, the authors conclude that the Wiener–Hopf approach is a very precise and efficient method for pricing options in the presence of multiple jumps.

The hedging of longer-term energy cash flows when many of the required hedging products are non-traded is the important practical topic of Chapter 29, by Leoni, Vandaele and Vanmaele. The authors point out that much attention in the literature has been focussed on risk-neutral spot price models for gas and electricity that immediately model short-to-delivery forwards and futures contracts, such as day-ahead contracts, under the assumption of perfect hedging. However, as we saw in Chapter 25, these markets do not trade these short-to-delivery contracts very far ahead in real time; in fact most practical forward trading concerns calendar years or seasons. This chapter treats the dynamics of such liquid longer-term forward contracts for these markets and shows how to hedge derivatives with smaller delivery periods, such as days or months, by applying a coarse-grained trading approach. Of course this practical approach introduces basis risk as the claims to be hedged are written on a different asset than those of the hedges, but in energy markets either the illiquidity does not allow direct hedging or the associated costs are too high. The prices of delivered energy products like electricity and gas are stated in currency per unit volume per unit of time leading to handling the forward price curve for energy flow by bucketting it into different time granularities. At the long end of the curve, far ahead in the future, forwards are for delivery of power over a calendar year, e.g., a CAL-17 contract prices delivery over all of 2017. As this year approaches, this year-long contract breaks down into quarterly contracts, which closer to delivery themselves break down into monthly, weekly and finally daily forward contracts. The chapter discusses electricity markets with base, peak and off-peak long forward contracts, separately from gas markets in which long forwards are broken down into delivery in winter or summer. The interest in option products in the European gas market is increasing with the most currently popular product being in a seasonal option, which is a strip of six monthly options that determine the pricing and hedging of such options. However, these components can only be hedged with seasonal forwards, providing a practical example of the problem the authors address in this chapter. They prove the existence of a locally risk-minimizing (LRM) strategy for payment streams in this setting making use of previous work in the literature initiated by Föllmer and Schweizer (1990). Noting that in very liquid markets such as oil markets where one has an established options market available, volatility skew and term structure arise which can be explained by a stochastic volatility factor. The first example the authors treat is the LRM in a Heston model. The second example is applicable to hedging seasonal gas options with season forward contracts and hedging peak and off-peak electricity options with base contracts. The authors provide numerical results comparing the rollover hedging costs of the near contract of their LRM strategy to the inaccurate delta hedging widely used in the power industry. They introduce an adjusted LRM strategy to take into account the volatilities and correlation of the two nearest contracts, which outperforms current industry hedging practice for both the GBM and the variance gamma Lévy, forward price processes.

The final two chapters in this volume point to critical future considerations for the power industry. In Chapter 30, Carmona, Coulon and Schwarz discuss the valuation of clean spread options, which link electricity, emissions and fuels. Building on their own recent previous work, and that of their colleagues, for this purpose, the authors construct a *structural* model. They assert that such a model is necessary to capture the economic factors entering power production, which are not captured by alternative, usually partial, reduced-form models. Specifically, they concentrate on electricity production and CO₂ emissions, defining a *clean spread option*, in the spirit of the more common *spark* (gas) and *dark* (coal) spread options, as the spread between the price at which 1 MWh of electricity can be sold in the market and the unit costs of production, including *both* gas (or coal) inputs and the market costs of carbon *allowances*. Market mechanisms for controlling CO₂ emissions by cap-and-trade schemes involving these traded allowances, whether they are voluntary or mandatory, have helped to price carbon emissions in Europe, the United States and globally. For the past decade, power generators (and other emitters) in Europe are given each year a number of allowances out of the total number for all industry, termed the *cap*, which must be surrendered at the end of the year against measured emissions. If their allowances are insufficient, a firm must go to the EU Emission Trading Scheme (ETS) and purchase allowances from firms in surplus, speculators and possibly investors. The ETS involves both a spot market and a futures market which, as is usual for other futures markets, leads spot allowance price discovery and can be used to anticipate future over- and under-supply of allowances. The value of clean spread options involving these allowances is expressed here in a formulation in which the demand for power and input fuel prices are the only factors with exogenous stochastic dynamics. The spot prices for power are based on a structural model for the *bid-stack*, in which bids to a central electricity delivery operator from generators are taken in bid price order, lowest (and least carbon emitting) first. The emission allowances are derivatives on a non-traded underlying. Their prices are computed using a forward backward stochastic differential system solved by applying a finite difference scheme to a multidimensional partial differential equation. Monte Carlo techniques are used to price clean spread options on emission allowances. These methods are used by the authors to investigate four indicative case studies. First, they compare the prices of spark and dark spread options in a market with no emission regulation with those of a cap-and-trade system in which the cap is increasingly strict. Next they analyze the effects that different fuel-type merit schemes for production have on option prices, and in the third case study they favourably compare option prices obtained with their structural scheme with those obtained from two alternative reduced-form schemes. Finally, they favourably compare from a policy viewpoint a cap-and-trade scheme using their structural model with a *fixed* carbon tax. The authors note the importance of the currently increased global market volatility in the price factors used in their study for designers of the increasing number of carbon emission regulation schemes. They conclude that their four case studies address the important considerations needed to understand the complex joint dynamics of electricity, emissions and fuels.

Medina and Pardo, in the final chapter of the volume, Chapter 31, conduct an empirical study of the properties of the futures prices of EUAs traded on the EU ETS, with a view to finding out if their stylized properties can be considered similar to financial assets, like either stocks and bonds or commodity futures. The answer is preliminary to addressing the theme of Section III, whether or not EUAs in the world's most developed cap-and-trade market (and those of other schemes) can be considered a diversifying asset for global investors. The authors make use of previous work on the stylized properties of financial assets (Cont 2001) and futures (Gorton et al. 2013) using daily data on EAU futures from the European Climate Exchange. Given that Phase 1 of the EU cap-and-trade scheme from 2005 to 2007 was considered to be a pilot and learning scheme, they concentrate on 2008, 2009 and 2010 EUAs of Phase 2 (which were expected to be carried forward into Phase 3 starting in 2013) and use futures data on them from April 2005 to the December date of the allowance maturity for each year. They also use the Euro Stoxx 50 index and its futures and fixed income futures for short-, medium- and long-term bond indices from the EUREX market, and Brent oil futures from the International Petroleum Exchange. One-month Euribor is taken as the risk-free return over the data period, and inflation is measured by the European Harmonized Index of Consumer Prices. Computing return histograms and using well-known statistical

methods for EUA futures returns, including GARCH modelling for their volatility, the authors found evidence for intermittency, short-term predictability, long-term memory (Taylor effect), volatility clustering, asymmetric volatility and positive relationships between volatility and trading volume and open interest. As a result, they feel that temporal dependencies should be modelled using ARMA–GARCH techniques and caution that the persistence of heavy tails for conditional returns implies that any unconditional return distributions should be fat-tailed. Unlike commodity futures returns, but like those of stocks and bonds, the study found negative asymmetry, positive correlation with equity index returns and higher volatility during trading sessions than when the allowance exchange is closed. Like commodity futures, the properties of inflation hedging and positive return correlations with bonds were found. The authors conclude that EUAs do *not* behave like either common commodity futures or financial assets, but have selected features of both, to form a new asset class with implications for volatility modelling, hedging use, portfolio analysis and cointegration of the carbon market with other economic and financial variables.

Epilogue

We hope that the contents of this book will contribute to the contemporary heated debates about commodity markets amongst investors, practitioners, academics and policymakers. Recent regulation regarding commodity trading is currently reducing the involvement of investment banks in commodity markets. In the medium term this may be expected to decrease commodity investment return volatility. We feel that only a deeper understanding of commodity futures markets and related derivative products will lead to financial practitioners creating truly innovative commodity-based products which are beneficial to investors.

M.A.H. Dempster

*University of Cambridge
Cambridge, United Kingdom*

Ke Tang

*Tsinghua University
Beijing, China*

References

- Ang, A. and Piazzesi, M., A no-arbitrage vector autoregression of term structure dynamics with macro-economic and latent variables. *J. Monetary Econ.*, 2003, **50**, 745–787.
- Benth, F.E., Benth, J.S. and Koekebakker, S. *Stochastic Modelling of Electricity and Related Markets*, 2008 (World Scientific: Singapore).
- Benth, F.E., Kholodnyi, V.A. and Laurence, P., *Quantitative Energy Finance: Modelling, Pricing, and Hedging in Energy and Commodity Markets*, 2014 (Springer: Berlin).
- Black, F., The pricing of commodity contracts. *J. Financ. Econ.*, 1976, **3**, 167–179.
- Black, F. and Scholes, M., The pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**, 637–659.
- Brennan, M., The supply of storage. *Am. Econ. Rev.*, 1958, **48**, 50–72.
- Brennan, M.J., Schwarte, E.S. and Lagnado, R., Strategic asset allocation. *J. Econ. Dyn. Control*, 1997, **21**, 1377–1403.
- Carmona, R. and Touzi, N., Optimal multiple stopping and valuation of swing options. *Math. Finance*, 2008, **18**, 239–268.
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Chang, E., Returns to speculators and the theory of normal backwardation. *J. Finance*, 1985, **40**, 193–208.

- Clark, E., Lesourd, J.-B. and Thiéblemont, R., *International Commodity Markets: Physical and Derivative Markets*, 2001 (Wiley: New York).
- Cont, R., Empirical properties of asset returns: Stylized facts and statistical issues. *Quant. Finance*, 2001, **1**, 223–236.
- Cox, J.C. and Ross, S.A., The valuation of options for alternative stochastic processes. *J. Financ. Econ.*, 1976, **3**, 145–166.
- Dempster, M.A.H. and Leemans, V., An automated FX trading system using adaptive reinforcement learning. *Expert Syst. Appl.*, 2006, **30**, 543–552.
- Dempster, M.A.H. and Tang, K., Estimating exponential affine models with correlated measurement errors: Applications to fixed income and commodities. *J. Bank. Finance*, 2011, **35**, 639–652.
- Dempster, M.A.H., Evstigneev, I.V. and Schenk-Hoppé, K.R., Growing wealth with fixed mix strategies. In *The Kelly Capital Growth Investment Criterion*, edited by L.C. MacLean, E.O. Thorpe and W.T. Ziemba, pp. 427–458, 2011 (World Scientific: Singapore).
- Föllmer, H. and Schweizer, M., Hedging of contingent claims under incomplete information. *Stochastics Mono.*, 1990, **5**, 389–414.
- Geman, H., *Commodities and Commodity Derivatives: Modelling and Pricing for Agriculturals, Metals and Energy*, 2005 (Wiley: New York).
- Gibson, R. and Schwartz, E., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Gorton, G., Hayashi, F. and Rouwenhorst, K., The fundamentals of commodity futures returns. *Rev. Finance*, 2013, **17**, 35–105.
- Henderson, B., Pearson, N. and Wang, L., New evidence on the financialization of commodity markets. *Rev. Financ. Studies*, 2015, **28**, 1285–1311.
- Jaillet, P., Ronn, E.I. and Tompaidis, S., Valuation of commodity-based swing options. *Manage. Sci.*, 2004, **50**, 909–921.
- Kaldor, N., Speculation and economic stability. *Rev. Econ. Studies*, 1939, **7**, 1–27.
- Kendall, M.G. and Bradford Hill, A., The analysis of economic time-series—Part I: Prices. *J. R. Stat. Soc. A*, 1953, **116**, 11–34.
- Keynes, J.M., Some aspects of commodity markets. *Manchester Guardian Comm.*, 1923, **13**, 784–786.
- Merton, R.C., Theory of rational option pricing. *Bell J. Econ. Manage. Sci.*, 1973, **4**, 141–183.
- Pirrong, C., *Commodity Price Dynamics: A Structural Approach*, 2012 (Cambridge University Press: Cambridge).
- Rockwell, C. Normal backwardation, forecasting and the returns to commodity futures traders. *Food Res. Inst. Studies*, 1967, **7**, 107–130.
- Routledge, B., Seppi, D. and Spatt, C., Equilibrium forward curves for commodities. *J. Finance*, 2000, **55**, 1297–1338.
- Samuelson, P.A., Proof that properly anticipated prices fluctuate randomly. *Industr. Manage. Rev.*, 1965, **6**, 41–43.
- Schwartz, E., The stochastic behaviour of commodity prices: Implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Tang, K., Time-varying long-run mean of commodity prices and the modelling of futures term structures. *Quant. Finance*, 2012, **12**, 781–790.
- Tang, K. and Xiong, W., Index investment and financialization of commodities. *Financ. Analyst. J.*, 2012, **68**, 54–74.
- Till, H. and Eagleeye, J., *Intelligent Commodity Investing: New Strategies and Practical Insights for Informed Decision Making*, 2007 (Risk Books: London).
- Wilford, D.S., Risk management insights from Markowitz optimization for constructing portfolios with commodity futures. *J. Financ. Perspect.*, 2014, **2**, 141–150.
- Working, H., A theory of the price of storage. *Am. Econ. Rev.*, 1949, **39**, 1254–1262.



Oil Products

1	Inconvenience Yield, or the Theory of Normal Contango	<i>Ilia Bouchouev</i>	3
2	Determinants of Oil Futures Prices and Convenience Yields	<i>M. A. H. Dempster, Elena Medova, and Ke Tang</i>	9
	Introduction • Oil Price Features • Three-Factor Model Statement • Model Interpretation • Conclusion		
3	Pricing and Hedging of Long-Term Futures and Forward Contracts with a Three-Factor Model	<i>Kenichiro Shiraya and Akihiko Takahashi</i>	31
	Introduction • Model • Estimation of Parameters • Futures Hedging Techniques • Stability of the Delta Hedge • Measuring the Distribution of Hedging Error Rates • Conclusion		
4	An Empirical Study of the Impact of Skewness and Kurtosis on Hedging Decisions	<i>Jing-Yi Lai</i>	53
	Introduction • The Hedging Model • The Econometrics Model • Empirical Results • Conclusion		
5	Long Term Spread Option Valuation and Hedging	<i>M. A. H. Dempster, Elena Medova, and Ke Tang</i>	71
	Introduction • Cointegrated Prices and Mean Reversion of the Spread • Modelling the Spread Process • Spread Option Pricing and Hedging • Crack Spread: Heating Oil/ WTI Crude Oil • Location Spread: Brent/WTI Crude Oil • Conclusion		
6	Analysing the Dynamics of the Refining Margin: Implications for Valuation and Hedging	<i>Andrés García Mirantes, Javier Población, and Gregorio Serna</i>	95
	Introduction • Data and Preliminary Findings • Common Long-Term Trend Factor Models • Crack-Spread Option Valuation • Conclusion		
7	Quantitative Spread Trading on Crude Oil and Refined Products Markets	<i>Mark Cummins and Andrea Bucca</i>	119
	Introduction • Crude Oil and Refined Products Markets • Optimal Statistical Arbitrage Trading Model • Multiple Hypothesis Testing: Data Snooping Bias • Outline of Empirical Analysis • Empirical Results • Conclusion		

Inconvenience Yield, or the Theory of Normal Contango

Ilia Bouchouev

Acknowledgements	8
References.....	8

The relative impact of fundamental imbalances and financial flows on oil prices remains the subject of intense debate among market participants, economists, and policy makers. Proponents of traditional fundamental views blame the oil price spike of 2008—when U.S. crude rose to an all-time high of almost \$150 a barrel—primarily on the inability of the global refining industry to cope with the rapid demand growth from emerging markets. Conversely, practitioners tend to point towards additional demand for futures from so-called ‘long-only’ investors, such as pension funds, whose increased allocations to commodities shifted established balances between buyers and sellers. Unfortunately, the economic objectives of financial investments in commodities are often poorly understood, leading to incorrect characterisations of the relevant market participants and misleading conclusions about the impact of their activities on the price dynamics.

This article highlights structural changes in supply and demand for hedging services which are becoming increasingly important factors for determining the price of oil. Traditionally, there has always been more producer hedging than consumer hedging in oil markets. While investor money used to balance this gap nicely, it now far exceeds it. It is important, though, to realise that the reasons investors are coming to the oil market have changed. With contango being the ‘normal’ oil market structure for the last several years, investors are no longer investing in the front end of the oil market primarily in order to collect a positive risk premium driven by the roll yield. Rather, they are hedging against inflation, U.S. dollar weakness, and possible geopolitical events that could negatively impact the rest of their investment portfolios. As the traditional roles of different players change, so too do the market instruments available to them. This article applies these changes in the behaviour of market participants to the classical Keynes–Hicks theory of normal backwardation, and the Kaldor–Working–Brennan theory of storage, and looks at how calendar spread options (CSOs) are becoming an increasingly popular risk management tool.

Traditional economics postulates that the price of commodities such as oil is determined by the equilibrium between consumer demand for physical barrels and supply of these barrels by producers. As both consumers and producers actively hedge their risks using financial instruments in order to reduce uncertainty, demand and supply for ‘paper’ barrels is created. Additional external capital is then required to absorb volatility and smooth price fluctuations caused by temporary supply and demand imbalances. This capital is more efficiently provided by investors via the futures and swap market to avoid extra costs associated with handling physical oil.



© 2012 iStockphoto LP

Physical and futures markets are interlinked, and for the purpose of price formation the supply and demand for both physical and paper barrels must be combined. The link between the two is provided by oil storage operators. When the demand from futures buyers exceeds supply causing futures to trade at a relative premium, the storage operator satisfies extra demand by selling futures and buying spot barrels from oil producers. In the spot market, the actions of the storage owner are identical to those of a consumer buying barrels, with a direct impact on supply and demand. The economic function of the storage operator is to convert extra financial demand for oil into physical demand when the spot prices are below futures, the state of the market defined as *contango*. On the other hand, when the spot market trades at a premium to futures—referred to as *backwardation*—the storage operator remains idle.

The study of the interplay between commodity producers, consumers, and risk absorbers goes back to Keynes (1930) and Hicks (1939). They assumed that producer hedgers exposed to larger and less diversified risks must dominate consumer hedgers, whose exposure to commodity prices is less concentrated. The Keynes–Hicks ‘theory of normal backwardation’ suggests that the futures curve must be downward sloping as commodity producers have to sell futures at a discount to spot prices in order to incentivise external capital to take on the producers’ risk of falling prices. In other words, the net hedging pressure in commodity markets must be negative. For many years, the commodity markets have been confirming this theory with backwardation prevailing and creating incentives for investors to buy futures at a discount and capture the risk premium given away by producers.

Historically, oil has been among the most backwardated commodities, attracting consistent attention from financial investors. The expected futures roll up was even granted a new term, the *roll yield*. It is well documented that a significant portion of the positive returns to early investors in oil futures came in the form of the roll yield rather than from the price appreciation. In addition, oil proved to be an excellent diversifier for investors because it had a low correlation to other financial classes and acted as a hedge against inflation, dollar weakness, and unpredictable geopolitical events. A good collection of articles describing commodities as a new asset class for investors is provided by Till and Eagleeye (2007).

Going back to historical analyses, Kaldor (1939) and Working (1949) interpreted the risk premium resulting from backwardation as a benefit to the owner of physical barrels, which they called the *convenience benefit*. Their equation also included the net storage and financing costs as follows:

$$\text{Spot Price} - \text{Future Price} = \text{Convenience Benefits} - \text{Storage/Financing Costs.}$$

However, as Kaldor himself pointed out, the equation may not hold when the markets are in steep *contango*, characterised by high inventories, negligible convenience benefits, and forward hedgers turning into net buyers. Such *contango* has been the consistent feature of the crude oil market in the U.S.

since roughly 2005. In fact, the change in the net hedging pressure from negative to positive around 2007 became one of the main stimuli for the accumulation of high inventories.

Estimating supply and demand for financial barrels and net hedging pressure is not straightforward. The vast majority of corporate hedging transactions are executed via the over-the-counter market. For that reason, any attempt to analyse the impact of hedgers and other financial market participants on oil prices based on exchange futures data is likely to be inconclusive and misleading. In an effort to gain some clarity with respect to the influence of hedging and financial flows, we assembled a comprehensive database of over-the-counter hedging volumes, utilising corporate reports and U.S. Securities and Exchange Commission (SEC) filings for major U.S. and Canadian corporate hedgers. In order to better understand the role of investors, Commodity Futures Trading Commission (CFTC) data aggregated since December 2007 on swap index positions was used. Prior to December 2007, the data was calculated using proprietary models and various interviews with market participants and index providers. Hedging and investment products linked to U.S. refined products were also included as they are typically converted to crude oil by refinery hedgers.

The results of this analysis, summarised in Figure 1.1, confirm Keynes' hypothesis that the majority of hedging volumes are executed by producers which results in negative hedging pressure. The hypothesis breaks down, however, if we add an additional demand for 'paper' barrels from passive investors. These financial investors have outpaced their original function of providing capital to balance the gaps between producers and consumers. In fact, capital provided by financial investors began to exceed the entire hedging imbalance it was supposed to stabilise around 2007. The net hedging pressure switched from negative to positive, from net sellers to net buyers.

With forward prices now systematically exceeding spot prices, any forward seller is now able to collect the new risk premium which is paid by investors. Figure 1.2 shows the magnitude of this risk premium. It presents cumulative returns from a passive investment in prompt WTI futures contracts at the beginning of each calendar year with futures rolled on the fifth business day of the month. It also breaks up such returns into the spot return, and the roll yield, which became increasingly negative. For example, in 2009 an investor in such a strategy would have experienced 67% negative roll yield which represented the real cost of holding an investment in oil futures. With this turn of events, it might be appropriate now to rename Kaldor's classical term into 'inconvenience yield', and the Keynes–Hicks' framework into 'the theory of normal contango'.

The answer can actually be explained by Keynes, Hicks and Kaldor themselves if we allow for some generalisation of their theories. It is still true that the provider of capital must require compensation in the form of a risk premium for transferring unwanted risks from the hedger. What has changed, though, is that financial market participants who were once classified as risk absorbers have become




FIGURE 1.1 Net hedging pressure.




FIGURE 1.2 WTI index performance.

net hedgers themselves. They are obviously no longer buying futures to collect the roll yield, as the Keynes–Hicks theory implied, because, as noted above, the roll yield has become a significant liability, to the extent that in some years they are losing money on their investment (see Figure 1.2). Instead, they are buying oil futures mostly as a hedge against inflation, and as protection for their larger investments in other asset classes. Whilst many investor products now exist which invest further along the oil curve to avoid the negative roll yield, the lion's share is still invested in the front month via indices such as the Standard & Poor's Goldman Sachs Commodity Index and Dow Jones–UBS Commodity Index. So, with the producer and the investor trading places, the structure of the oil futures market has changed.

What remains to be determined is who now collects the risk premium paid by macro-economic hedgers in the form of negative roll yield. As we saw earlier, the hedging by oil producers was not sufficient to meet growing demand from investors. Moreover, if upcoming derivatives regulations impose stricter margin requirements for over-the-counter oil transactions, the credit and liquidity capacity required by independent producers to support large-scale hedging programmes will likely decline even further.

Recall the function of the storage operator in the contango market who sells relatively expensive futures to financial hedgers and buys relatively cheap physical oil in the spot market from producers. It is the oil warehouse who now collects the premium (i.e. the negative roll yield) for taking extra risks. Going back to Kaldor's equation, we can propose a modified version to capture this risk premium:

$$\text{Spot Price} - \text{Futures Price} = \text{Convenience Benefits} - \text{Storage/Financing Cost} - \text{Risk Premium.}$$

Brennan (1958) suggested a very similar equation but in a slightly different context. He associated the risk premium with an additional risk held by producers for carrying unhedged inventories. Nowadays, most of the barrels that go into commercial storage are hedged with futures, and the risk premium has acquired a new meaning. It has become equivalent to the return required by the storage operator for taking risks of investing in new storage capacity in order to meet the growing demand of the new hedger, the financial investor.

Despite attractive short-term economics during contango markets, building new storage requires significant capital investment, yet such investments cannot promise long-term profits. If the hedging demand from such financial investors wanes, the market could return to backwardation with no payoff accrued to the storage owner. The cash flows from owning a storage asset are identical to owning a put option on the difference between spot and future prices. Rearranging the previous equation, and taking

only the positive values (since an asset owner retains the option not to run its business when profit margins are negative), we can express this option as follows:

$$\text{Risk Premium} = \text{Max}(\text{Strike Price} - (\text{Spot Price} - \text{Future Price}), 0),$$

where

$$\text{Strike Price} = \text{Convenience Benefits} - \text{Storage/Financing Costs}.$$

In other words, the risk premium—or, alternatively, the value of the storage asset—is the market value of the put option on the futures term-structure which, in turn, is determined by the expected volatility of future spreads. The time value of this option is linked to the risk premium that financial oil investors should be paying to operators to incentivise them to construct new storage facilities, the availability of which will flatten contango and reduce the negative roll yield.

Such financial options, known in the market as calendar spread options (CSOs), have been trading over-the-counter for over a decade. Selling these options and collecting risk premiums allows storage operators to raise money for building new storage facilities. Derivatives dealers, who also assume a certain amount of credit risk, market these options to speculators, and more recently investor hedgers. For financial investors, buying such options became an attractive alternative to paying the negative roll yield. In 2006 these options were listed at NYMEX/CME, and they recently became one of the fastest growing derivatives instruments in the oil market (see Figure 1.3).

It is possible to estimate the fair value of storage by using the market price of tradable CSOs. The details of such valuation are left for future publications. Overall, investors would have been much better off hedging their negative yield by buying CSOs. An investment in a CSO is economically equivalent to an ownership interest in a physical asset. The cost of such investment is significantly cheaper than paying the negative roll yield to a warehouse, in the form of a storage fee. Financial demand from investors continues to incentivise construction of new storage facilities, which mitigates the risks of oil price volatility for U.S. consumers.




FIGURE 1.3 CME WTI CSO open interest.

This article discusses important changes in the behaviour of oil market participants and explains the resulting impact on the relationship between spot and futures prices. The risk premium that a producer used to pay to the provider of the risk capital in the backwardated market has transformed into the risk premium paid by a financial investor to a storage operator in a contango market. This risk premium which could be estimated using over-the-counter option prices is typically smaller than the cost to the financial investor to roll the futures. The option prices also indicate the market-based fair value of the storage. The oil price is now influenced by the demand from financial investors and the supply of futures from storage operators, and not only by traditional balances between producers and consumers.

Acknowledgements

Various aspects of this paper were presented at Global Derivatives (May 2010 and April 2011, Paris), Energy Risk Asia (September 2010, Singapore), and Energy Risk USA (May 2011, Houston). The author thanks participants of these conferences for useful discussions. The information herein is taken from sources the author believes to be reliable. All of the contents are subjective opinions of the author alone.

References

- Brennan, M., The supply of storage. *Am. Econ. Rev.*, 1958, **47**(1), 50–72.
Hicks, J.R., *Value and Capital*, 1939 (Oxford University Press: Oxford).
Kaldor, N., Speculation and economic stability. *Rev. Econ. Stud.*, 1939, **7**(1), 1–27.
Keynes, J., *A Treatise on Money*, Vol. II, 1930 (Macmillan: London).
Till, H. and Eagleeye, J., *Intelligent Commodity Investment*, 2007 (Risk Books: New York).
Working, H., The theory of price of storage. *Am. Econ. Rev.*, 1949, **39**(December), 1254–1262.

Determinants of Oil Futures Prices and Convenience Yields

2.1	Introduction	9
2.2	Oil Price Features	11
	Term Structure of Oil Futures Open Interest • Single-Factor Convenience Yield Models	
2.3	Three-Factor Model Statement	15
	Dynamics of Spot Prices • Two-Factor Convenience Yield • Results Comparison with Two Factor Models	
2.4	Model Interpretation.....	20
	Explanatory Variable Specification • SVAR Model Statement • Results	
2.5	Conclusion	26
	Acknowledgements	26
	References.....	26
	Appendix A	28

M. A. H. Dempster

Elena Medova

Ke Tang

Commodity futures prices are usually modelled using affine term structure spot price models with latent factors extracted from the data. However, very little research to date has considered the question—What are the economic drivers behind the calibrated latent factors? This paper addresses this question in the context of a three-factor—short-, medium- and long-term—model for crude oil spot prices by studying the relations between these factors and appropriate economic variables. An affine combination of the short- and medium-term factors is identified as the (instantaneous) convenience yield. Estimating a structural vector auto-regression model we find that the short-term factor mainly relates to demand variables in the physical markets and to trading variables in the futures markets (such as the net short position of commercial hedgers), the medium-term factor relates to business cycles, demand and trading variables, and the long-term factor relates mainly to financial factors.

Keywords: Oil futures; Futures term structure; Theory of storage; Theory of normal backwardation; Kalman filter; SVAR

2.1 Introduction

Commodity futures exchanges arose in the late 19th and early 20th centuries to allow the forward contract mitigation of cyclical supply/demand imbalances between agricultural producers and consumers while limiting speculation through requiring the posting of margin on futures positions. In the 21st

century, expanded to a wide range of commodities and related services and increasingly electronic, they continue to serve their original purpose. But they are also vital to the forecasting of spot prices by large resource producers for management evaluation of project alternatives and investment opportunities over very long-term horizons. Global resource firms may or may not hedge their physical activities in the futures markets, but increasingly they have come to see that sophisticated price forecasting is a prerequisite to the use of real option techniques for forward planning and risk management of ongoing operations.

However, for use in exploration, acquisition evaluation, or project development and risk management, senior management cannot be content with reduced form 'black box' price forecasting methods devoid of an economic understanding of the commodity markets involved. Focussing on crude oil prices, this paper attempts to meet these stringent managerial criteria by specifying a three-factor spot price model for oil and studying its relationship to different economic variables, which includes financial variables (such as SP500 returns, US dollar returns, etc.), business cycle variables (such as the business cycle coincident index), demand variables (such as inventory and the heating-crude oil spread) and trading variables (such as futures open interest growth and hedging pressure). We hope to contribute to an understanding of the relationships between oil prices, physical inventory management, financial hedging and speculation. Although this paper treats the oil markets, the model treated here may be applied to a wider range of commodities, upon which we are currently engaged (see e.g. Dempster and Tang 2011).

Commodities are real assets and so their prices should be influenced by their supply and demand. The *theory of storage* (Kaldor 1939, Working 1949, Brennan 1958) sees optimal inventory management as the main determinant of commodity prices. But since commodities are traded through futures contracts, financial markets and trading behaviour will also influence the commodity prices and term structure, as noted by Keynes (1930) in his *theory of normal backwardation*. More recently, Bailey and Chan (1993) showed that financial factors such as the spread between BAA and AAA bonds can influence the convenience yield of many commodities. In this paper, we shall examine the impact both of supply, demand and business variables and of financial and trading factors on the movement and shape of the oil futures term structure.

Both producers and consumers wish to make forecasts for long-term planning and investment decisions since commodity prices represent their output revenues and input costs, respectively. Various authors have expressed differing opinions on the long-term evolution of commodity prices. Most (see, e.g., Cuddington and Urzua [1987] and Gersovitz and Paxson [1990]) believe that commodity prices are non-stationary since it is hard statistically to reject the most parsimonious geometric random walk model using historical time series data. Cashin *et al.* (2000) have shown that shocks to commodity prices are typically long lasting, while Grilli and Yang (1988) found that real primary commodity prices have a trend of about 0.5% a year using a dataset from 1900 to 1986. Schwartz and Smith (2000) use *geometric Brownian motion* (GBM) to model such long-term behaviour because of the ability of GBM to capture trend and the persistency of shocks. However, Bessembinder *et al.* (1995) discovered strong mean reversion in commodity log spot prices, suggesting that a *geometric Ornstein Uhlenbeck* (GOU) process might be more appropriate. Although not pointed out explicitly, Casassus and Collin-Dufresne (2005) use a mean-reverting process to model log spot commodity prices. Geman and Nguyen (2005) use a mean-reverting log spot price with stochastic mean and stochastic volatility to model soybean futures prices.* Many economists believe that commodity prices in the medium term are closely related to the business cycle (e.g., Fama and French 1988), which is usually considered to be a mean-reverting process. Short-term swings in commodity prices have substantial impacts for many speculators and short-term strategic investors and the short-term factors driving commodity prices are usually also considered to be mean-reverting (e.g., Schwartz and Smith 2000).

* Using soybean inventory data, Geman and Nguyen (2005) also show that soybean futures return volatility is negatively related to soybean inventory (or positively related to 'scarcity', the reciprocal of inventory), which is consistent with the theory of storage.

In this paper, we use a mean-reverting process to model the short-term factor in our three-factor (log) spot price model. However, it is not appropriate to model short-term influences using only a *single* mean-reverting factor, since one factor alone cannot adequately model the complicated short- to medium-term behaviour of commodity prices. This suggests that two mean-reverting factors—one short- and one medium-term—are needed to model price movements.* Our medium-term factor captures business cycles and long-term demand in the global economy, while a long-term GBM factor captures trend-related persistent shocks, such as technology growth, long-term supply through the discovery of new resources, etc. Intuitively, the time scale of the short-term factor is several months, that of the medium-term factor 1 to 2 years and that of the long-term factor decades or even longer. The three-factor model treated here nests several other models, including those of Gibson and Schwartz (1990) and Schwartz and Smith (2000), which are equivalent. Because of oil's importance to the global economy, and the ease of obtaining inventory data for it and relevant economic variables, we use crude oil futures to illustrate our model's development and to examine the various intuitions presented briefly above.

After developing the model† in state space form, we use the Kalman filter to obtain the estimated historical paths of the three latent factors from observations of oil futures prices. We then perform a *structural vector auto-regression* (SVAR) analysis involving the three estimated factor paths and the historical paths of several economic variables, including financial, business-cycle, fundamental and trading variables. We find that financial variables mainly affect the long-term p factor and the business cycle variable influences mainly the medium-term y factor. The demand variables affect both x and y factors, and higher net demand results in both a higher short-term x factor (deeper short-end backwardation) and a higher y factor (deeper long-end backwardation). The trading variables also influence both x and y factors, and more intensive trading and stronger hedging pressure result in higher factor levels.

The paper is organized as follows. Section 2.2 examines several features of WTI crude oil prices and develops a detailed motivation for a three-factor spot price model. Section 2.3 presents such a model and explains its relationship to earlier models. Section 2.4 examines the relationship between the three estimated latent factors and several economic variables. Section 2.5 concludes.

2.2 Oil Price Features

In this section we characterize features of crude oil futures prices and their evolution.

2.2.1 Term Structure of Oil Futures Open Interest

Weekly oil futures prices and open interest for WTI crude oil (CL) traded on the New York Mercantile Exchange (NYMEX) were obtained from 1986.06 to 2010.12 from Pinnacle Data Corp. The times to maturity of these futures contracts range from several days to about 17 months* (the first to the seventeenth contract). Figure 2.1 shows futures term structure shapes commonly observed in the market with their observation dates. At the long end we can see both *contango* (top two diagrams) and *backwardation* (bottom two diagrams) term structures, while at the short end we see U and hump shapes. Thus the short end of the term structure appears to be more volatile and is not necessarily conformal with the long end.

Futures prices are discovered through trading, so to investigate their short-term behaviour in Figure 2.2 we plot the futures *open interest*, i.e. the total number of futures contracts that have not expired, or been fulfilled by delivery, against their times to maturity. The figure demonstrates that the main oil trading activities are concentrated in the one month futures contract. Its open interest is much

* The medium-term factor should obviously revert to its mean more slowly than the short-term factor.

† It is an affine futures term structure model for the log spot price in the $A_0(3)$ form of Dai and Singleton (2000).

‡ The open interest for futures with time to maturity longer than 17 months is very small (see Figure 2.4). Also, in the earlier part of our dataset, futures prices with maturity longer than this are not available.




FIGURE 2.1 The oil futures term structure.




FIGURE 2.2 Average open interest vs. time to maturity.

larger than the value obtained from an exponential function fit to open interest at other maturities.* This large open interest in the nearby contract indicates a uniquely high liquidity which can result in different behaviour of short-term oil futures prices. Furthermore, investors and hedgers all prefer to use short-term futures instead of long-term futures and thus contribute to this behaviour. For example, passive commodity index investors tend to invest in short-term rather than long-term futures (for details, refer to Tang and Xiong 2010). As shown by Culp and Miller (1995), Mellon and Parsons (1995) and Brennan and Crew (1997), hedgers employ short-term futures to hedge longer-term obligations.

2.2.2 Single-factor Convenience Yield Models

In previous research, nearly all researchers have used a single factor to model *convenience yield*, which is the commodity equivalent of equity dividend yield representing the opportunity return to physical ownership of the commodity. Among such models, the Gibson–Schwartz (1990) two-factor model, the Schwartz–Smith (2000) two-factor model and the Schwartz three-factor model are commonly used. We use the Schwartz–Smith model here to investigate whether or not one factor is enough to model convenience yields.

2.2.2.1 Short Term Pricing Errors

Table 2.1 contains the log pricing errors from Schwartz (1997, p. 939) and Schwartz–Smith (2000, p. 903), where the short-term (1 month) futures prices have a noticeably larger error than the others.[†] This is consistent with the hypothesis that short-term futures price movements are often not conformal with longer-term futures price movements and thus a *separate* factor is needed to model short-term futures price movements.

2.2.2.2 Convenience Yields Inferred from the Schwartz–Smith (2000) Model

In the Gibson–Schwartz model, log futures prices are expressed as an affine combination of two latent factors: the log spot price and the convenience yield. Thus given parameter estimates for the model the convenience yield and spot price can be backed out using any two futures prices. We re-estimate the Gibson–Schwartz model using our own dataset and in Figure 2.3 plot the evolution of the model convenience yields inferred from 1 and 3 month futures and 15 and 17 month futures and their difference. It is clear that these two estimates of convenience yield are *strongly inconsistent*. The unconditional standard deviation of the convenience yield implied from 1 and 3 month futures is 26.3% per annum, while that from 15 and 17 months is 34.7% p.a. The unconditional standard deviation of the difference between these two convenience yields is surprisingly large, about 33.6% p.a.* The differences between these two convenience yield maturities should therefore not be overlooked, but should instead be modelled using a new short-term factor.

TABLE 2.1 Model Errors for Futures with Different Times to Maturity

Maturity (Months)	Gibson–Schwartz Two-Factor Model	Schwartz–Smith Two-Factor Model	Schwartz Three- Factor Model
1	0.043	0.042	0.045
5	0.006	0.006	0.007
9	0.003	0.003	0.003
13	0	0	0
17	0.004	0.004	0.004

* We fit the open interest frequency by fitting $P_t = a \exp(b\tau)$ to it, where P_t is the (time) average open interest for futures contracts with time to maturity.

† When we estimate the Schwartz–Smith (2000) model using our dataset we find a similar phenomenon.

* Note that the implied convenience yield during the financial crisis fluctuates widely.




FIGURE 2.3 Implied convenience yields and their difference for the two-factor Gibson and Schwartz model.




FIGURE 2.4 Estimated latent factor evolution in the three-factor model.

TABLE 2.2 The Principal Components of the Implied Convenience Yield

Component	Variance Explained (%)
First	92.05
Second	5.99
Third	1.33
Fourth	0.39
Fifth	0.22
Sixth	0.02

2.2.2.3 Principal Component Analysis on Convenience Yields

As we have seen, one factor is not enough to model the convenience yield; in this section we test how many factors are actually needed using a *principal component analysis* (PCA). Since convenience yield is not directly observable, we infer the implied convenience yield $\delta(t, T)$ from commodity futures prices and interest rates using

$$\begin{aligned}\delta(t, T) &= r_t - \frac{\ln(F(t, T)) - \ln S_t}{T - t} \\ &\approx r_t - \frac{\ln(F(t, T)) - \ln(F(t, T_0))}{T - T_0},\end{aligned}\quad (2.1)$$

where T_0 is the time to maturity of nearby futures contracts, T is the time to maturity of futures with a relatively longer horizon and r_t is the instantaneous rate corresponding to the three month LIBOR rate. We use nearby and 3, 6, 9, 12, 15, 17 month futures (i.e. $T = 3, 6, 9, 12, 15, 17$) to calculate six time series of implied convenience yield over our data period and then perform a PCA on them. Table 2.2 shows the variance explained by each factor. Clearly, two factors can explain more than 98% of the overall variance of the convenience yield, so that a two-factor model is good enough to catch its behaviour.

In the sequel we use two factors to model convenience yield, one short term and one for a longer medium-term horizon. The short-term factor should correct both the large pricing error of short-term contracts in two-factor models and the mismatch arising with these models of implied convenience yields which are backed out from short- and longer-maturity futures. Adding a long-term factor, the resulting three-factor model can capture the different shapes of the futures term structure shown in [Figure 2.2](#).

2.3 Three-factor Model Statement

We begin by modelling the dynamics of the log spot oil price $\mathbf{G}$ in terms of convenience yield using two factors.

2.3.1 Dynamics of Spot Prices

In the *market* (physical) *measure* the system is given by,*

$$d\mathbf{G}_t = \left(r^f + \lambda_G - \delta_t - \gamma_t - \frac{1}{2}\sigma_G^2 \right) dt + \sigma_G d\mathbf{W}_G, \quad (2.2)$$

$$d\delta_t = k_\delta(\alpha - \delta_t)dt + \sigma_\delta d\mathbf{W}_\delta, \quad (2.3)$$

$$d\gamma_t = -k_\gamma\gamma_t dt + \sigma_\gamma d\mathbf{W}_\gamma, \quad (2.4)$$

$$Ed\mathbf{W}_G d\mathbf{W}_\delta = \rho_{\delta G} dt, Ed\mathbf{W}_\delta d\mathbf{W}_\gamma = \rho_{\delta \gamma} dt, Ed\mathbf{W}_G d\mathbf{W}_\gamma = \rho_{G\gamma} dt. \quad (2.5)$$

* Boldface is used throughout to denote random entities, here conditional.

Here, at time t , $\mathbf{G}_t := \ln(\mathbf{S}_t)$ is the logarithm of the spot price, $\delta_t + \gamma_t$ is the spot (instantaneous) convenience yield with the medium-term δ_t and the short-term γ_t mean-reverting factors having long-run means α and 0 respectively in the market measure, λ_G is the *market price of risk* premium of the G process and $\mathbf{W}_G$, $\mathbf{W}_\delta$ and $\mathbf{W}_\gamma$ are Wiener processes with σ_G , σ_δ and σ_γ their corresponding volatilities.

In the *risk-neutral measure* this system becomes

$$d\mathbf{G}_t = \left(r^f - \delta_t - \gamma_t - \frac{1}{2}\sigma_G^2 \right) dt + \sigma_G d\mathbf{W}_G^Q, \quad (2.6)$$

$$d\delta_t = k_\delta (\alpha - \delta_t - \lambda_\delta) dt + \sigma_\delta d\mathbf{W}_\delta^Q, \quad (2.7)$$

$$d\gamma_t = k_\gamma (-\gamma_t - \lambda_\gamma) dt + \sigma_\gamma d\mathbf{W}_\gamma^Q, \quad (2.8)$$

$$\begin{aligned} Ed\mathbf{W}_G^Q d\mathbf{W}_\delta^Q &= \rho_{\delta G} dt, \quad Ed\mathbf{W}_\delta^Q d\mathbf{W}_\gamma^Q = \rho_{\delta\gamma} dt, \\ Ed\mathbf{W}_G^Q d\mathbf{W}_\gamma^Q &= \rho_{G\gamma} dt, \end{aligned} \quad (2.9)$$

where $k_\delta \lambda_\delta$ and $k_\gamma \lambda_\gamma$ are, respectively, the market risk premia for the δ and γ processes.

Setting the γ factor identically equal to zero, (2.2) to (2.9) becomes the Gibson–Schwartz (1990) model so that our model is its extension, but with convenience yield decomposed into two parts, δ and γ , with different mean-reversion speeds.

Defining $x_t := (1/k_\delta)(\delta_t - \alpha)$, $y_t := \gamma_t/k_\gamma$ and $\mathbf{p}_t := \mathbf{G}_t - \mathbf{x}_t - \mathbf{y}_t$ in the market measure, we have

$$d\mathbf{x}_t = \frac{1}{k_\delta} d\delta_t = -k_\delta x_t dt + \frac{\sigma_\delta}{k_\delta} d\mathbf{W}_\delta, \quad (2.10)$$

$$d\mathbf{y}_t = \frac{1}{k_\gamma} d\gamma_t = -k_\gamma y_t dt + \frac{\sigma_\gamma}{k_\gamma} d\mathbf{W}_\gamma, \quad (2.11)$$

$$\begin{aligned} d\mathbf{p}_t &= d\mathbf{G}_t - \frac{d\delta_t}{k_\delta} - \frac{d\gamma_t}{k_\gamma} = \left(r^f + \lambda_G - \alpha - \frac{1}{2}\sigma_G^2 \right) dt \\ &\quad + \sigma_G d\mathbf{W}_G - \frac{\sigma_\delta}{k_\delta} d\mathbf{W}_\delta - \frac{\sigma_\gamma}{k_\gamma} d\mathbf{W}_\gamma. \end{aligned} \quad (2.12)$$

Setting $k_x := k_\delta$, $k_y := k_\gamma$, $\sigma_x := \sigma_\delta/k_\delta$, $\sigma_y := \sigma_\gamma/k_\gamma$, $\lambda_x := \lambda_\delta/k_\delta$, $\lambda_y := \lambda_\gamma/k_\gamma$, $\lambda_p := \lambda_G - \lambda_\delta - \lambda_\gamma$, $u := r^f + \lambda_G - \alpha - (1/2)\sigma_G^2$, $\sigma_p^2 := \sigma_G^2 + \sigma_x^2 + \sigma_y^2 + 2\rho_{\delta\gamma}\sigma_x\sigma_y - 2\rho_{G\delta}\sigma_G\sigma_x - 2\rho_{G\gamma}\sigma_G\sigma_y$ and $d\mathbf{W}_x := d\mathbf{W}_\delta$, $d\mathbf{W}_y := d\mathbf{W}_\gamma$, $d\mathbf{W}_p := (1/\sigma_p)(\sigma_G d\mathbf{W}_G - (\sigma_\delta/k_\delta) d\mathbf{W}_\delta - (\sigma_\gamma/k_\gamma) d\mathbf{W}_\gamma)$, the original model in the market measure becomes

$$\ln(\mathbf{S}_t) = \mathbf{x}_t + \mathbf{y}_t + \mathbf{p}_t, \quad (2.13)$$

$$d\mathbf{x}_t = -k_x x_t dt + \sigma_x d\mathbf{W}_x, \quad (2.14)$$

$$d\mathbf{y}_t = -k_y y_t dt + \sigma_y d\mathbf{W}_y, \quad (2.15)$$

$$d\mathbf{p}_t = u dt + \sigma_p d\mathbf{W}_p, \quad (2.16)$$

$$\begin{aligned} Ed\mathbf{W}_x d\mathbf{W}_y &= \rho_{xy} dt, \quad Ed\mathbf{W}_x d\mathbf{W}_p = \rho_{xp} dt, \\ Ed\mathbf{W}_y d\mathbf{W}_p &= \rho_{yp} dt, \end{aligned} \quad (2.17)$$

where $\mathbf{x}$ is the *short-term* factor with mean-reversion speed k_x and volatility σ_x , $\mathbf{y}$ is the *medium-term* factor with mean-reversion speed k_y and volatility σ_y , $\mathbf{p}$ is the *long-term trend* factor with growth rate u and volatility σ_p and W_x, W_y and W_p are all Wiener processes. Note that factors $\mathbf{x}$ and $\mathbf{y}$ both have zero long-run means so that they will fluctuate around the trend factor $\mathbf{p}$.

In the risk-neutral measure the above system becomes

$$\ln(\mathbf{S}_t) = \mathbf{x}_t + \mathbf{y}_t + \mathbf{p}_t, \quad (2.18)$$

$$d\mathbf{x}_t = k_x(-\mathbf{x}_t - \lambda_x)dt + \sigma_x d\mathbf{W}_x^Q, \quad (2.19)$$

$$d\mathbf{y}_t = k_y(-\mathbf{y}_t - \lambda_y)dt + \sigma_y d\mathbf{W}_y^Q, \quad (2.20)$$

$$d\mathbf{p}_t = (u - \lambda_p)dt + \sigma_p d\mathbf{W}_p^Q \quad (2.21)$$

$$\begin{aligned} Ed\mathbf{W}_x^Q d\mathbf{W}_y^Q &= \rho_{xy} dt, \quad Ed\mathbf{W}_x^Q d\mathbf{W}_p^Q = \rho_{xp} dt, \\ Ed\mathbf{W}_y^Q d\mathbf{W}_p^Q &= \rho_{yp} dt, \end{aligned} \quad (2.22)$$

where $k_x \lambda_x$, $k_y \lambda_y$ and λ_p are the *risk premia* of factors $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{p}$ respectively.*

We term the model (2.13) to (2.22) the *three factor (log) spot price* model. It determines spot price fluctuations in terms of three components: two mean-reverting factors representing short- and medium-term economic forces and one long-term factor that reflects the equilibrium commodity price trend and captures permanent price shocks. We note that the model belongs to the exponential affine class in the framework of Duffie *et al.* (2000).†

Solving (2.19), (2.20) and (2.21), and substituting into (2.18), together with taking logarithms of the no-arbitrage condition for the price at t of the futures contract with maturity T given by

$$F(t, T) = E_t^Q[\mathbf{S}_T] \quad (2.23)$$

in terms of the conditional expectation in the risk-neutral measure Q at t , yields $\ln F(t, T)$ in terms of the three factors at t as

$$\begin{aligned} \ln F(t, T) &= (x_t + \lambda_x) e^{-k_x(T-t)} + (y_t + \lambda_y) e^{-k_y(T-t)} \\ &\quad + p_t - (\lambda_x + \lambda_y) + (u - \lambda_p)(T-t) \\ &\quad + \frac{1}{2} \left[\begin{aligned} &\frac{1 - e^{-2k_x(T-t)}}{2k_x} \sigma_x^2 \\ &+ \frac{1 - e^{-2k_y(T-t)}}{2k_y} \sigma_y^2 + \sigma_p^2 (T-t) \\ &+ \frac{2(1 - e^{-(k_y + k_x)(T-t)})}{k_x + k_y} \rho_{xy} \sigma_x \sigma_y \\ &+ \frac{2(1 - e^{-k_x(T-t)})}{k_x} \rho_{xy} \sigma_x \sigma_p \\ &+ \frac{2(1 - e^{-k_y(T-t)})}{k_y} \rho_{yp} \sigma_y \sigma_p \end{aligned} \right]. \end{aligned} \quad (2.24)$$

* Note that, since these risk premia for the $\mathbf{x}$ and $\mathbf{y}$ factors are *constant*, the mean-reversion speeds of these factors are the same under the market and risk-neutral measures (see Casassus and Collin-Dufresne [2005] who assume risk premia stochastic).

† See last footnote of Section 2.1.

2.3.2 Two-factor Convenience Yield

The price at t of the futures contract with maturity T is given in terms of the instantaneous convenience yield at t by

$$F(t, T) = S_t \exp \left(r^f (T - t) - \int_t^T \delta(t, s) ds \right), \quad (2.25)$$

where $\delta(t, s)$ is the *instantaneous convenience yield* at time t of the contract with maturity s . In our model

$$\begin{aligned} \delta(t, T) = & r^f + (\mathbf{x}_t + \lambda_x) k_x e^{-k_x(T-t)} + (\mathbf{y}_t + \lambda_y) k_y e^{-k_y(T-t)} \\ & - (u - \lambda_p) - \frac{1}{2} \left[e^{-2k_x(T-t)} \sigma_x^2 + e^{-2k_y(T-t)} \sigma_y^2 + \sigma_p^2 \right. \\ & + 2\rho_{xy} \sigma_x \sigma_y e^{-(k_x + k_y)(T-t)} + 2\rho_{xp} \sigma_x \sigma_p e^{-k_x(T-t)} \\ & \left. + 2\rho_{yp} \sigma_y \sigma_p e^{-k_y(T-t)} \right]. \end{aligned} \quad (2.26)$$

When $T \rightarrow t$ this reduces to the (*instantaneous*) *spot convenience yield* given by

$$\begin{aligned} \delta_t = \delta(t, t) = & r^f + k_x (\mathbf{x}_t + \lambda_x) + k_y (\mathbf{y}_t + \lambda_y) - u + \lambda_p \\ & - \frac{1}{2} (\sigma_x^2 + \sigma_y^2 + \sigma_p^2 + 2\rho_{xy} \sigma_x \sigma_y \\ & + 2\rho_{xp} \sigma_x \sigma_p + 2\rho_{yp} \sigma_y \sigma_p), \end{aligned} \quad (2.27)$$

so that the spot convenience yield δ is seen to be an affine combination of $\mathbf{x}$ and $\mathbf{y}$ factors as designed. Our calibration of the three-factor model for oil futures shows, as expected, that the $\mathbf{x}$ factor has a much higher mean-reversion speed than that of the $\mathbf{y}$ factor (see Table 2.3). As a consequence, (2.25) and (2.26)

TABLE 2.3 Parameter Estimates of Two- and Three-Factor Models

Variable	Three-Factor Model	Two-Factor Model
k_x	3.4152 (0.1147)	
k_y	0.8802 (0.0384)	1.0715 (0.0166)
u	0.0809 (0.0425)	0.0838 (0.0428)
σ_x	0.1977 (0.0078)	
σ_y	0.2817 (0.0078)	0.2866 (0.0067)
σ_p	0.1953 (0.0051)	0.1960 (0.0044)
λ_x	-0.0205 (0.0128)	
λ_y	0.1600 (0.0701)	0.0720 (0.0581)
λ_p	0.0731 (0.0425)	0.1063 (0.0428)
ρ_{xy}	-0.0794 (0.0518)	
ρ_{xp}	0.0838 (0.0431)	
ρ_{yp}	-0.0067 (0.0472)	0.0932 (0.0346)
ξ_1	0.0160 (0.0037)	0.0347 (0.0072)
ξ_2	0.0004 (0.0000)	0.0115 (0.0025)
ξ_3	0.0014 (0.0006)	0.0000 (0.0000)
ξ_4	0.0016 (0.0005)	0.0022 (0.0006)
ξ_5	0.0011 (0.0007)	0.0021 (0.0009)
ξ_6	0.0158 (0.0033)	0.0161 (0.0034)
ξ_7	0.0160 (0.0035)	0.0172 (0.0036)
Log-likelihood	24493	22331

Note: $\xi_1, \xi_2, \xi_3, \xi_4, \xi_5, \xi_6, \xi_7$ are, respectively, the mean absolute pricing errors of the F1, F3, F6, F9, F12, F15 and F17 contracts. The quantities in parentheses are (asymptotic) standard deviations.

imply that the longer-term convenience yields $\delta(t, T)$ are determined mainly by the $\mathbf{y}$ factor, while the spot convenience yield δ_t is determined mainly by the $\mathbf{x}$ factor.

Although convenience yield is a concept of the theory of storage, in the context of the theory of normal backwardation we wish to know the shape and overall slope (i.e. *contango*, upwards slope, or *backwardation*, downwards slope) of the futures price term structure. Taking logarithms of both sides of (2.25) and differentiating the result with respect to maturity gives

$$\frac{\partial(F(t, T))/\partial T}{F(t, T)} = r^f - \delta(t, T). \quad (2.28)$$

The convenience yield $\delta(t, T)$ therefore determines the sign of the slope of the term structure of futures prices and the $\mathbf{x}$ and $\mathbf{y}$ factors can be regarded as two components of this slope. Thus, when the instantaneous convenience yield $\delta(t, T)$ is strictly less than the instantaneous risk-free rate r^f , futures prices are in *contango locally* in maturity. However, because $\mathbf{x}$ and $\mathbf{y}$ may have different signs at a specific time t , the three-factor model is capable of reproducing the empirical near term U or humped futures curves of [Figure 2.1](#).

2.3.3 Results

The appendix shows the state space form of the three-factor model needed for the parameter estimation filtering technique.* Since the three factors of our model are not directly observable, to calibrate it we use the *EM algorithm* procedure which alternates between the *Kalman filter* and maximum likelihood parameter estimation of the model in state space form to convergence (Schwartz 1997, Schwartz and Smith 2000, Geman and Nguyen 2005).

By F1, F3, F6, F9, F12, F15, and F17 we denote respectively the 1st, 3rd, 6th, 9th, 12th, 15th and 17th month futures contracts (in the order of their maturities) which we use in the calibration of the three-factor spot price model. [Table 2.3](#) shows the parameter estimates. These estimated parameters are nearly all significant except for the risk premia. Both the $\mathbf{x}$ and $\mathbf{y}$ factors are significantly mean-reverting, as can be seen in [Figure 2.4](#), which shows the estimated paths of the three factors. From the estimated parameters the short-term factor $\mathbf{x}$ has a half-life of about 2.5 months with a volatility of 20%, the medium-term factor $\mathbf{y}$ has a half-life of about 9.5 months with volatility 28% and the long-term factor $\mathbf{p}$ has a volatility of about 20%.

2.3.4 Comparison with Two Factor Models

To see whether or not the three-factor model is significantly better than two-factor models, we remove the $\mathbf{x}$ factor and re-estimate the resulting model with only the $\mathbf{y}$ and $\mathbf{p}$ factors.† [Table 2.3](#) demonstrates that the pricing errors ($\xi_1, \dots, \xi_7$) are generally smaller than those reported for the Gibson–Schwartz model by Schwartz (1997). Comparing our two- and three-factor models, the inclusion of the $\mathbf{x}$ factor significantly improves the data fit, according to the likelihood ratio test,‡ and reduces pricing errors for short-term contracts.

* Note that in order to make a comparison with the results of the two-factor models of Gibson and Schwartz (1990) and Schwartz and Smith (2000), we employ the same estimation method used in those papers. We also used the method proposed by Dempster and Tang (2011) to eliminate mean-reversion parameter estimation errors, but the results reported here are little changed.

† We have seen in [Section 2.3.1](#) that this two-factor model is the same as the Gibson-Schwartz (1990) model.

‡ Note that this test applies to our nested case. The likelihood ratio test statistic is 4324, highly significant at the 99% confidence level for the Chi squared distribution with five degrees of freedom.

2.4 Model Interpretation

In this section we study the relations between the historical paths of various economic variables and those of our latent factors (Figure 2.4), estimated in terms of the means of the sequential posterior Gaussian state distributions obtained from the Kalman filter using the optimal parameter estimates of the final iteration of the EM algorithm. Since commodities are both real and financial assets, we expect that both fundamental and financial variables will play an important role in explaining the three factors.

2.4.1 Explanatory Variable Specification

We classify our explanatory variables into four categories: (1) variables from other financial markets, such as US dollar index returns, SP500 equity index returns, etc.; (2) variables indicating the phase of the business cycle, such as the coincident business cycle index and the term spread on US interest rates; (3) variables indicating net demand for oil, such as the oil inventory level and the heating oil–crude oil spread; (4) trading variables, such as growth rate of open interest and hedging pressure for oil futures contracts.

More specifically, we utilized the following variables at weekly frequency.

2.4.1.1 Financial Variables

- **USD (y1):** Weekly returns of the US dollar index. This index measures the performance of the US dollar against a basket of currencies. It goes up when the US dollar gains strength relative to other currencies.
- **SP500 (y2):** Weekly returns of the S&P500 equity index. Inclusion of SP500 returns controls for the possibility that investors were pursuing trading strategies in oil futures that are conditional on equity markets.
- **LIBOR (y3):** The weekly level of the three month US LIBOR rate. Casassus and Colin-Dufresne (2005) and Frankel (2008) show that interest rates tend to influence the willingness to store inventory and thus will influence the convenience yield.*
- **VIX (y4):** The weekly level of the VIX index for the equity market. This index represents one measure of the market's expectation of stock market volatility over the next 30-day period. It is a weighted blend of prices for a range of options on the S&P 500 index.
- **CreditSpread (y5):** The weekly spread between Moody's BAA-rated and AAA-rated corporate bond yields, following Bailey and Chan (1993), as a proxy for the default premium.

2.4.1.2 Business Cycle Variables

- **TermSpread (y6):** The weekly spread between the 10 year and three month Treasury bond yields. The term spread between the long-term and short-term interest rates has often been found to be the most important predictor of economic recessions (see, e.g., Estrella and Mishkin [1998]).
- **CoinIndex (y7):** The *change* in weekly levels of the business cycle coincident index, as this index, obtained from the Conference Board Inc., is not stationary. However, since the coincident index is calculated at monthly frequency, we assume the weekly values of coincident index change are identical in each month.

* For example, a higher interest rate corresponds to a higher marginal cost of storage, a higher convenience yield and a futures term structure more likely to be in backwardation, as Keynes (1930) described. The current low interest rates have led through the opposite effect to an oil futures market in contango and a frenzy of oil storage building to exploit it physically (Bouchouev 2011).

2.4.1.3 Demand Variables

- **HOCLSpread (y8):** The weekly *log* spread between heating and crude oil prices. Since heating oil is the main product of crude oil, the (log) price spread can reflect the relative scarcity of crude oil (Casassus *et al.* 2010). Note that a high spread means that crude oil is cheap relative to heating oil and hence has low convenience yield, a phenomenon independent of heating oil, and hence spread, seasonality.
- **Inventory (y9):** US weekly crude oil inventory (excluding strategic petroleum reserves) in millions of barrels obtained from the US Energy Information Administration. Since this data is non-stationary, we follow Gorton *et al.* (2007) in first applying a Hodrick–Prescott (HP) filter to the whole time series. The detrended stationary part is used in our empirical analysis.

2.4.1.4 Trading Variables

- **OpenInterest (y10):** The weekly growth rate (log difference) of open interest for all oil contracts, obtained from the US Commodity Futures Trading Commission (CFTC).
- **HedgingPressure (y11):** The weekly hedging pressure, obtained from the US Commodity Futures Trading Commission (CFTC). This measure is calculated using the hedgers' short position less their long position normalized by total open interest.

2.4.2 SVAR Model Statement

We estimate a *structural vector auto-regression* (SVAR) model (Sims 1980, Hamilton 1994) to address the relationship between the three latent factors and the 11 explanatory variables, namely

$$BY_t = AY_{t-1} + \Sigma \varepsilon_Y, t, \quad (2.29)$$

where the vector of the variables is given by $Y_t := (y_{1,t}, \dots, y_{11,t}, z_t')'$, with z_t a 3 vector representing the latent **x**, **y** and **p** factors, A and B are 14×14 matrices and $\varepsilon_{Y,t}$ is a vector of Gaussian disturbances with a spherical covariance matrix.

In structuring the SVAR matrix B we assume that all variables can influence the latent **x**, **y** or **p** factors, but that, consistent with our three-factor model, these factors do not (directly) influence each other. Similarly, the financial and business cycle variables do not directly influence each other, but they do affect the fundamental and trading variables. On the other hand, these latter variables do *not* influence financial and business cycle variables, i.e. the two blocks of variables are in cascade (causal Wold recursive) form. We also assume that the heating oil–crude oil spread can influence oil inventory and *vice versa*, i.e. they are determined simultaneously. Further, we assume that the fundamental demand and supply variables can influence futures trading and *vice versa*, but that open interest and hedgers' positions do not influence each other. These assumptions lead to a B matrix of (2.29), corresponding to each of the three latent factors *separately* in turn, in lower triangular form given by

$$B = \begin{bmatrix} 1 & & & & & & & & & & & & & & \\ 0 & 1 & & & & & & & & & & & & & \\ 0 & 0 & 1 & & & & & & & & & & & & \\ 0 & 0 & 0 & 1 & & & & & & 0 & & & & & \\ 0 & 0 & 0 & 0 & 1 & & & & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 1 & & & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & & & & & & & & \\ x & x & x & x & x & x & x & 1 & & & & & & & \\ x & x & x & x & x & x & x & x & 1 & & & & & & \\ x & x & x & x & x & x & x & x & x & 1 & & & & & \\ x & x & x & x & x & x & x & x & x & 0 & 1 & & & & \\ x & x & x & x & x & x & x & x & x & x & 1 & & & & \end{bmatrix},$$

with x representing non-zero elements. Note that, by using this B matrix, the full model can be estimated under our assumptions by separately estimating the resulting model for each latent factor in turn.

2.4.3 Results

To see how the exogenous variables influence our three latent variables, we first analyse the impulse response functions of the estimated SVAR model. Figures 2.5, 2.6 and 2.7 show, for each of the x , y and p factors respectively, impulse response functions for a one standard deviation positive shock to each of the exogenous explanatory variables in turn. Note that the vertical scales in these diagrams are variable. The gray area in each represents the 95% confidence level obtained from bootstrapping.

2.4.3.1 x Factor Impulse Responses

The LIBOR rate impacts positively on the short-term convenience yield x factor, which is consistent with the standard argument of the theory of storage, i.e. a higher interest rate will result in a higher marginal cost of storing commodities. Thus a higher LIBOR rate should correspond to a higher convenience yield. We also see that credit spread co-moves with the x factor, which is consistent with Acharya *et al.* (2008) in that a higher default risk tends to encourage more commodity producers' hedging and thus a futures curve in deeper backwardation.




FIGURE 2.5 Impulse responses of the x factor to a one standard deviation positive shock of each of the 11 explanatory variables.




FIGURE 2.6 Impulse responses of the y factor to a one standard deviation positive shock of each of the 11 explanatory variables.

The log spread between heating and crude oil influences the x factor negatively. A high spread means a relatively low price of crude oil and less demand for it (or more crude oil inventory), corresponding to a lower convenience yield. For details of the equilibrium relationship between oil convenience yield and the heating oil-crude oil spread, see Casassus *et al.* (2010).

As shown by Hong and Yogo (2010), open interest in commodity futures forecasts commodity returns. Our result shows that a higher open interest growth rate corresponds to a futures term structure in deeper backwardation. We also see that more hedging pressure corresponds to deeper backwardation. This is consistent with Keynes (1930) and Hirshleifer (1990), i.e. the more futures contracts sold the deeper the backwardation of the futures term structure.

2.4.3.2 y Factor Impulse Responses

Similar to the situation with the x factor, we see that the LIBOR rate also impacts positively on the medium-term convenience yield y factor, which is again consistent with the theory of storage. We also see that the VIX correlates positively with the y factor. Since the VIX index can be seen as an indicator of the volatility of financial markets, the theory of storage says that higher volatility will lead to larger convenience yields, resulting here in a positive relationship between the VIX and the y factor.




FIGURE 2.7 Impulse responses of the p factor to a one standard deviation positive shock of each of the 11 explanatory variables.

Both the term spread and the business cycle coincident index have a positive impact on the y factor. This is because a high term spread and coincident index both correspond to a booming state of the economy with a higher oil demand and hence a higher convenience yield. The log spread between heating and crude oil also has a negative impact on the y factor; the explanation for this is similar to that for the x factor. The inventory has a positive impact on the y factor, which is consistent with the theory of storage, i.e. high inventory results in higher convenience yield and deeper backwardation of the futures term structure (see (2.28)). Similar to the x factor, the y factor is also affected by the growth of open interest and the hedging pressure, but at lower impact levels.

2.4.3.3 p Factor Impulse Responses

First observe from Figure 2.7 that the impulse responses of the long-term p factor to explanatory variable shocks are of larger magnitude and converge faster to equilibrium than those of the x and y factors, as is consistent with its GBM dynamics.

The US dollar index co-moves negatively with the permanent shock p factor. This is because oil is traded in dollars, hence the depreciation of dollars should increase the price of oil due to the numeraire effect. This effect tends to be ‘permanent’, i.e. only affecting the long-term p factor in our model.

TABLE 2.4 Forecast Error Variance Decomposition

Variable	x Factor		y Factor		p Factor	
	1-Week Horizon	13-Week Horizon	1-Week Horizon	13-Week Horizon	1-Week Horizon	13-Week Horizon
USD	0.02	0.2	0.03	0.06	2.93	3.51
SP500	0.00	0.25	0.00	0.35	2.60	3.14
LIBOR	0.62	0.70	0.48	0.53	0.61	0.60
VIX	0.15	0.65	0.43	0.12	0.82	0.88
Credit spread	0.38	0.17	0.26	0.08	0.23	0.23
Term spread	0.00	0.50	1.21	0.80	0.02	0.07
Coin index	0.01	0.14	0.01	1.85	0.00	0.00
HO—CO spread	11.88	4.84	1.03	8.31	0.64	1.59
Inventory	0.89	0.03	0.08	2.11	0.00	0.27
Open interest	1.68	0.55	0.80	0.60	0.83	0.79
Hedging pressure	3.00	2.38	11.16	12.24	0.59	0.70
<i>Aggregate averages</i>						
Financial	0.23	0.39	0.24	0.23	1.44	1.67
Business cycle	0.01	0.32	0.61	1.33	0.01	0.04
Fundamental	6.39	2.44	0.56	5.21	0.32	0.93
Trading	2.34	1.47	5.98	6.42	0.71	0.75

Note: This table reports the forecasting error variance decomposition for each factor over two rolling forecast horizons—one week and one quarter. Quantities in the table are percentages of the total factor variance.

There have been several studies of the relationship between the stock market portfolio and futures prices (e.g., Dusak [1973] and Holthausen and Hughes [1978]). Using *t*-tests they found that no correlation exists between futures and market portfolio returns. However, by decomposing oil futures prices into three factors, we see that the long-term **p** factor does co-move with SP500 index returns, but the **x** and **y** factors do not.

Similar to the **x** and **y** factors, the heating oil–crude oil spread has a negative impact on the **p** factor as well.

2.4.3.4 Forecast Error Variance Decomposition

We complement the conclusions derived from impulse response analysis of our estimated SVAR model with the forecast error variance decomposition from one- and 13-step ahead rolling forecasts (see, e.g., Lütkepohl [2007]). Table 2.4 shows the variance decomposition for both one week and one quarter forecasts, with entries that are percentages of the forecast error variance of each factor accounted for by exogenous shocks to each explanatory variable, or an average of these percentages for a group of variables.

We see that the variability of the log spread between heating and crude oil strongly influences that of the **x** factor; however, this effect decreases rapidly as the forecasting horizon lengthens. Trading variables such as the growth of the open interest and the hedging pressure also play an important role. The **y** factor is influenced by the variability of the business cycle variables, term spread and coincident index. The fundamental demand and supply variables, heating-crude oil spread and the inventory also play a role, but their role is only significant at the 13-week forecasting horizon. Futures hedging pressure has a very strong influence on the variability of the medium-term **y** factor, but the long-term **p** factor is mainly influenced by the US dollar index and SP500 returns.

2.4.3.5 Summary

The financial variables mainly affect the $\mathbf{p}$ factor, however they also have a minimal influence on the $\mathbf{x}$ and $\mathbf{y}$ factors. The business cycle variables influence mainly the $\mathbf{y}$ factor; a booming state corresponding to larger $\mathbf{y}$ factors (deeper long-end backwardation). The fundamental supply and demand variables affect both $\mathbf{x}$ and $\mathbf{y}$ factors; higher net demand results in higher $\mathbf{x}$ factor (deeper short-end backwardation) and $\mathbf{y}$ factor (deeper long-end backwardation) levels. The trading variables influence both $\mathbf{x}$ and $\mathbf{y}$ factors; more intensive futures trading and stronger hedging pressure result in higher $\mathbf{x}$ and $\mathbf{y}$ factors.

2.5 Conclusion

In this paper we find that the two-factor models in the literature are not able to model the whole crude oil futures price term structure, especially at the short end. Hence we propose a three-factor model for commodity futures prices. This model is shown to be an extension of the Gibson–Schwartz (1990) (Schwartz–Smith 2000) model.

An affine combination of the $\mathbf{x}$ and $\mathbf{y}$ factors in our model represents convenience yield, while the third $\mathbf{p}$ factor models long-term trend. By regressing the three factors on several economic variables using an SVAR model, we see that the short-term $\mathbf{x}$ factor is highly correlated with demand and trading variables. The medium-term $\mathbf{y}$ factor has a relationship with the business cycle, net oil demand and trading variables. The long-term $\mathbf{p}$ factor is mainly related to financial variables. The business cycle and fundamental variables affect the movement and the shape of the oil futures price term structure; but financial and trading variables do as well. This phenomenon reflects the fact that commodities combine the characteristics of both real and financial assets.

Acknowledgements

We thank participants at the Cambridge Finance Seminar, Global Derivatives, Trading and Risk Management and the annual Cambridge-Princeton Conference for comments, Professors R. Carmona, A. Harvey, M. Pesaran and L.C.G. Rogers for helpful discussions, and three anonymous referees for criticism that forced us to materially improve the paper. Tang acknowledges financial support from the National Natural Science Foundation of China (grant No. 71171194).

References

- Acharya, V., Lochstoer, L. and Ramadorai, T., Does hedging affect commodity prices? The role of producer default risk. Working paper, London Business School, 2008.
- Bailey, W. and Chan, K., Macroeconomic influences and the variability of the commodity futures basis. *J. Finance*, 1993, **48**, 555–573.
- Bessembinder, H., Coughenour, J., Seguin, P. and Smoller, M., Mean-reversion in equilibrium asset prices: Evidence from the futures term structure. *J. Finance*, 1995, **50**, 361–375.
- Bouchouev, I., The inconvenience yield or the theory of normal contango. *Energy Risk*, September 2011.
- Brennan, M., The supply of storage. *Am. Econ. Rev.*, 1958, **48**, 50–72.
- Brennan, M. and Crew, N., Hedging long maturity commodity commitments with short-dated futures contracts. In *Mathematics of Derivative Securities*, edited by M.A.H. Dempster and S.R. Pliska, 1997 (Cambridge University Press: New York).
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Casassus, J., Liu, P. and Tang, K., Long-term economic relationships and correlation structure in commodity markets. Working Paper, Cornell University, 2010.

- Cashin, P., Liang, H. and McDermott, J., How persistent are shocks to world commodity prices. *IMF Staff Papers*, 2000, **47**, 177–217.
- Cuddington, J. and Urzu, M., Trends and cycles in primary commodity prices. Working Paper, Georgetown University, Washington, DC, 1987.
- Culp, C. and Miller, M., Metallgesellschaft and the economics of synthetic storage. *J. Appl. Corp. Finance*, 1995, **7**, 62–76.
- Dai, Q. and Singleton, K., Specification analysis of affine term structure models. *J. Finance*, 2000, **55**, 1943–1978.
- Dempster, M.A.H. and Tang, K., Estimating exponential affine models with correlated measurement errors: Applications to fixed income and commodities. *J. Bank. Finance*, 2011, **35**, 639–652.
- Duffie, D., Pan, J. and Singleton, K., Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 2000, **68**, 1343–1376.
- Dusak, K., Futures trading and investor returns: An empirical investigation of commodity risk premium. *J. Polit. Econ.*, 1973, **81**, 1387–1406.
- Estrella, A. and Mishkin, F., Predicting U.S. recessions: Financial variables as leading indicators. *Rev. Econ. Statist.*, 1998, **80**, 45–61.
- Fama, E. and French, K., Business cycles and the behavior of metals prices. *J. Finance*, 1988, **43**, 1075–1093.
- Frankel, J., An explanation for soaring commodity prices. EUVOX, 25 March 2008.
- Geman, H. and Nguyen, V., Soybean inventory and forward curve dynamics. *Mgmt Sci.*, 2005, **51**, 1076–1091.
- Gersovitz, M. and Paxson, C., The economies of Africa and the prices of their exports. Working Paper, Princeton Studies in International Finance No. 68, Princeton University, 1990.
- Gibson, R. and Schwartz, E., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Gorton, G., Hayashi, F. and Rouwenhorst, K., The fundamentals of commodity futures returns. Yale ICF Working Paper No. 07-08, Yale University, 2007.
- Grilli, E. and Yang, M., Primary commodity prices, manufactured goods prices and the terms of trade of developing countries: What the long run shows. *World Bank Econ. Rev.*, 1988, **2**, 1–47.
- Hamilton, J., *Time Series Analysis*, 1994 (Princeton University Press: Princeton NJ).
- Hirshleifer, D., Hedging pressure and futures price movements in a general equilibrium model. *Econometrica*, 1990, **58**, 411–428.
- Hong, H. and Yogo, M., Commodity market interest and asset return predictability. Working Paper, Princeton University, 2010.
- Holthausen, D. and Hughes, J., Commodity returns and capital asset pricing. *Financ. Mgmt*, 1978, **7**, 37–44.
- Kaldor, N., Speculation and economic stability. *Rev. Econ. Stud.*, 1939, **7**, 1–27. Keynes, J.M., *A Treatise on Money*, 1930 (Macmillan: London).
- Lütkepohl, H., *Introduction to Multiple Time Series Analysis*, 2007 (Springer: Berlin).
- Mello, A. and Parsons, J., Maturity structure of a hedge matters: Lessons from the Metallgesellschaft debacles. *J. Appl. Corp. Finance*, 1995, **8**, 106–120.
- Schwartz, E.S., The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Schwartz, E. and Smith, J., Short term variations and long term dynamics in commodity prices. *Mgmt Sci.*, 2000, **46**, 893–911.
- Sims, C., Macroeconomics and reality. *Econometrica*, 1980, **48**, 1–48.
- Tang, K. and Xiong, W., Index investment and financialization of commodities. *Financial Analysts Journal*, 2012, **68**, 54–74.
- Working, H., The theory of the price of storage. *Am. Econ. Rev.*, 1949, **39**, 1254–1262.

Appendix A

Three-Factor Model in State Space Form

The state space form of a dynamic statistical model consists of a transition and a measurement equation. The *transition equation* describes the dynamics of the data-generating process of unobservable *state variables*. In our model this is a discrete-time version of (2.14) to (2.16). The *measurement equation* relates a multivariate time series of *observable variables*, here the future prices of different maturities, to the unobservable vector of state variables, the $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{p}$ factors. The measurement equation is obtained from (2.24) by adding *uncorrelated* noise to take into account *pricing errors*.* These errors may be caused by bid–ask spreads, non-simultaneity of the observations, etc.

In more detail, suppose the data are sampled at equally spaced times t_n , $n = 1, \dots, N$, and that $\Delta := t_{n+1} - t_n$ is the interval between two observations. Let $X_n := [x_{t_n} y_{t_n} p_{t_n}]'$ represent the vector of state variables at time t_n where the prime denotes transpose. Discretizing (2.14) to (2.16) we obtain the *transition equation* as

$$\mathbf{X}_{n+1} = A\mathbf{X}_n + b + \mathbf{w}, \quad (\text{A.1})$$

where $\mathbf{w}$ is a Gaussian random noise vector with mean 0 and covariance matrix Q and A , b and Q are given by

$$A = \begin{bmatrix} e^{-k_x \Delta} & 0 & 0 \\ 0 & e^{-k_y \Delta} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (\text{A.2})$$

$$b = [0 \ 0 \ u\Delta]', \quad (\text{A.3})$$

$$Q = \begin{bmatrix} \sigma_x^2 \frac{1 - e^{(-2k_x \Delta)}}{2k_x} & \frac{\rho_{xy} \sigma_x \sigma_y}{k_x + k_y} (1 - e^{-(k_x + k_y) \Delta}) \frac{\rho_{xp} \sigma_x \sigma_p}{k_x} (1 - e^{-k_x \Delta}) \\ \frac{\rho_{xp} \sigma_x \sigma_p}{k_x + k_y} (1 - e^{-(k_x + k_y) \Delta}) & \sigma_x^2 \frac{1 - e^{(-2k_y \Delta)}}{2k_y} & \frac{\rho_{yp} \sigma_y \sigma_p}{k_y} (1 - e^{-k_y \Delta}) \\ \frac{\rho_{xp} \sigma_x \sigma_p}{k_x} (1 - e^{-k_x \Delta}) & \frac{\rho_{yp} \sigma_y \sigma_p}{k_y} (1 - e^{-k_y \Delta}) & \sigma_p^2 \Delta \end{bmatrix} \quad (\text{A.4})$$

Let $Z_n := [\ln F(t, t + \tau_1), \dots, \ln F(t, t + \tau_M)]'$ represent log futures prices, where $\tau_1, \dots, \tau_M$ are the *times to maturity* for these 1, $\dots$, M futures contracts. From (2.24) the *measurement equation* becomes

$$Z_n = C_n X_n + d_n + \varepsilon_n \quad (\text{A.5})$$

where

$$C_n = \begin{bmatrix} e^{-k_x \tau_1} & e^{-k_y \tau_1} & 1 \\ \vdots & \vdots & \vdots \\ e^{-k_x \tau_M} & e^{-k_y \tau_M} & 1 \end{bmatrix}, \quad (\text{A.6})$$

* But see Dempster and Tang (2011) for more general assumptions.

$$d_n = \begin{bmatrix} \lambda_x (e^{-k_x \tau_1} - 1) + \lambda_y (e^{-k_y \tau_1} - 1) + (u - \lambda_p) \tau_1 \\ + \frac{1}{2} \left(\frac{1 - e^{-2k_x \tau_1}}{2k_x} \sigma_x^2 + \frac{1 - e^{-2k_y \tau_1}}{2k_y} \sigma_y^2 + \sigma_p^2 \tau_1 \right. \\ \left. + 2 \frac{1 - e^{-(k_x + k_y) \tau_1}}{k_x + k_y} \rho_{xy} \sigma_x \sigma_y + 2 \frac{1 - e^{-k_x \tau_1}}{k_x} \rho_{xp} \sigma_x \sigma_p + 2 \frac{1 - e^{-k_y \tau_1}}{k_y} \rho_{yp} \sigma_y \sigma_p \right) \\ \vdots \\ \lambda_x (e^{-k_x \tau_M} - 1) + \lambda_y (e^{-k_y \tau_M} - 1) + (u - \lambda_p) \tau_M \\ + \frac{1}{2} \left(\frac{1 - e^{-2k_x \tau_M}}{2k_x} \sigma_x^2 + \frac{1 - e^{-2k_y \tau_M}}{2k_y} \sigma_y^2 + \sigma_p^2 \tau_M \right. \\ \left. + 2 \frac{1 - e^{-(k_x + k_y) \tau_M}}{k_x + k_y} \rho_{xy} \sigma_x \sigma_y + 2 \frac{1 - e^{-k_x \tau_M}}{k_x} \rho_{xp} \sigma_x \sigma_p + 2 \frac{1 - e^{-k_y \tau_M}}{k_y} \rho_{yp} \sigma_y \sigma_p \right) \end{bmatrix}, \quad (\text{A.7})$$

and ε_n is an error term allowing noise in the sampling of data with covariance matrix*

$$H = \begin{bmatrix} \xi_1^2 & 0 & 0 \\ \vdots & \vdots & \vdots \\ 0 & 0 & \xi_M^2 \end{bmatrix}. \quad (\text{A.8})$$

* But see Dempster and Tang (2011) for more general assumptions.

3

Pricing and Hedging of Long-Term Futures and Forward Contracts with a Three-Factor Model

3.1	Introduction	32
3.2	Model	33
3.3	Estimation of Parameters	35
	Estimation Results • Comparison with Actual Data	
3.4	Futures Hedging Techniques	40
	Hedging Error Rate of the Parallel Hedge • Hedging Error Rate of the Delta Hedge	
3.5	Stability of the Delta Hedge	45
	Out of Sample Hedges • Hedging Long-term Forward Contracts	
3.6	Measuring the Distribution of Hedging Error Rates	48
	Distribution of Hedging Error Rates	
3.7	Conclusion	49
	Acknowledgements	50
	References	50
	Appendix A: Expectation and Covariance Matrix of State Variables	51
	Appendix B: The Number of Futures Units	52

Kenichiro Shiraya

Akihiko Takahashi

This paper demonstrates the pricing and hedging efficiency of a three-factor stochastic mean reversion Gaussian model of commodity prices using oil and copper futures and forward contracts. The model is estimated using NYMEX WTI (light sweet crude oil) and LME Copper futures prices and is shown to fit the data well. Furthermore, it shows how to hedge based on a three-factor model and confirms that using three different futures contracts to hedge long-term contracts outperforms the traditional parallel hedge based on a single futures position using time series data and simulation. It also finds that the three-factor model outperforms the two-factor version with respect to the replication of actual term structures and that stochastic mean reversion models outperform constant mean reversion models in Out of Sample hedges.

Keywords: Commodity prices; Multi-factor models; Hedging errors; Hedging techniques

3.1 Introduction

The term structure of commodity futures exhibits complex shape changes and a number of different models have been proposed for its estimation. In this paper, we propose a three-factor model to estimate the term structure of commodity futures, and then propose and verify effective hedging techniques for long-term futures and forwards estimated with the model, using as hedging instruments the short- and medium-term futures that are tradable.

Black (1976) advocated the idea of trading commodities as ‘equities without dividends’ and made use of geometric Brownian motion. However, given the complexity of the shapes associated with the term structure of commodities futures, the simple geometric Brownian motion model proposed by Black (1976) is not a good fit. To resolve this problem, mean reversion has been introduced. Unlike equities, when commodities prices rise, there is generally (albeit with a time lag) an increase in supply; conversely, when prices decline, supply decreases. The fact that prices are determined by the supply and demand balance means that the supply side adjusts supply volumes, which has the effect of constraining the potential for commodities prices to move in a single direction. That is why it is generally considered appropriate to employ mean reversion in commodity pricing models. Much empirical research has been carried out on this subject. For example, it was verified by Bessembinder *et al.* (1995).

Nonetheless, even if mean reversion is used in a one-factor model, it is difficult to represent the complex term structure of commodities futures, leading Gibson and Schwartz (1990) to propose a model that supplements the fluctuation of spot prices with a convenience yield stochastic process, and Schwartz (1997) to propose a model that explicitly employs convenience yields and interest rates as stochastic factors.

On the other hand, different methods have been proposed that do not attempt to individually model commodities’ spot prices, convenience yields or interest rates, but instead attempt direct modelling using state variables with the mean reversion of spot prices. Examples of direct modelling of spot prices include Schwartz and Smith’s (2000) two-factor mean reversion model, Casassus and Collin-Dufresne’s (2005) three-factor mean reversion model and Cortazar and Naranjo’s (2006) *N*-factor mean reversion model.

We use three-factor Gaussian models with or without constant mean reversion. The models’ parameters are estimated using a Kalman filter and have been confirmed to reproduce actual futures prices on the NYMEX WTI (light sweet crude oil) and LME Copper markets. Where our research differs from prior research is that we study cases both with and without a constant mean reversion level in the commodities price model and provide a detailed analysis not only of the model’s ability to reproduce futures prices, but also its utility in hedging.

Commodities hedging is a long-debated topic. For example, Culp and Miller (1995), Mello and Parsons (1995) and many others have discussed it in terms of the Metallgesellschaft case. Culp and Miller (1995) explain that, like equities, etc., the forward prices for commodities are determined by the mechanism of ‘cost of carry’ and argue that long-term forward contracts can be hedged by holding short-term futures and rolling over the contract months. On the other hand, Mello and Parsons (1995) acknowledge that it is possible to use short-term futures to hedge long-term forward contracts, but criticize the hedging technique employed by Metallgesellschaft, which was to use a unit of short-term futures to hedge a unit of long-term forward contracts. They use the Gibson and Schwartz (1990) model to demonstrate that short-term prices are more sensitive to spot price changes than long-term prices and that the actual number of short-term futures required to hedge one unit of long-term forward contracts is approximately 0.3. Because of this, the trading of Metallgesellschaft, while having hedging elements, is deemed to be primarily futures speculation. Schwartz (1997) also comments on this point, using one- to three-factor models to calculate hedge positions and explaining that when one factor is used, the position is significantly less than 1, approximately 0.2–0.4, and even with two and three factors it is still, on a net basis, less than 1. Neuberger (1999) uses multiple contracts to hedge long-term exposure and shows the benefits of the simultaneous use of different hedging instruments.

Examples of research analysing not only hedge positions but also hedging errors include Brenann and Crew (1997), Korn (2005) and Buhler *et al.* (2004). Brenann and Crew (1997) attempt to use a number

of different expiring futures as hedging instruments for hedges using a two-factor model, but all of the futures used as hedges expire within 6 months, and the futures to be hedged are also extremely short at no more than 2 years. Buhler *et al.* (2004) use several different models to compare and analyse performance when hedging 10-year forward contracts. However, the futures used as hedging instruments are extremely short, expiring in no more than 2 months, and the data are also only up to 1996, therefore, this analysis does not incorporate the rapid rises in commodities prices seen in recent years. Korn (2005) showed the hedging error with one- and two-factor models, but he did not show it with a three-factor model.

In this paper, we compare the hedging error using Metallgesellschaft's parallel hedging and multi-factor model based hedging. More specifically, we verify the stability of hedges based on two- and three-factor models that do and do not have a constant mean reversion level, and we provide a detailed analysis of the differences in hedge effectiveness due to differences in the way in which state variables are calculated and differences in the required futures units, and hedging error rate distribution (based on simulations) due to differences in the contract months of the futures used as hedging instruments. We also use time series data to verify hedges for long-term forward contracts, for which interest rate factors have been taken into account. We find that the three-factor model without constant mean reversion level can effectively hedge long-term futures against the complex changes in term structures of recent years.

In Section 3.2, we propose a three-factor model including a two-factor model as a special case, which does not explicitly incorporate interest rates or convenience yields, and we use that model to derive an analytic solution for futures prices. Section 3.3 makes use of Kalman filters to estimate the model's parameters. Section 3.4 makes use of short- and medium-term futures to create a hedging technique for long-term futures and to analyse performance when this hedging strategy is used. Section 3.5 takes a more practical approach, analysing hedges on 'Out of Sample' and long-term forward contracts. Section 3.6 uses a simulation to analyse how the form of the distribution changes for the hedge error rate depending upon the selection of futures contract months. In the appendix, we provide the expectation and covariance of the model expressed in futures prices and notes on the numbers of units of nearer maturity futures required to hedge long-term futures.

3.2 Model

We first describe a three-factor Gaussian model used for pricing and hedging futures and forward contracts. S_t represents the spot price of a commodity at time t . The logarithm of spot price at this time is expressed by the following equation:

$$\log S_t = x_t^1, \quad (3.1)$$

where x^1 expresses a state variable corresponding to the spot price of the commodity and follows the stochastic differential equation

$$\begin{aligned} dx_t^1 &= \kappa(x_t^2 + x_t^3 - x_t^1)dt + \sigma_1 dW_t^1, \\ dx_t^2 &= -\gamma x_t^2 dt + \sigma_2 dW_t^2, \\ dx_t^3 &= (\alpha - \beta x_t^3)dt + \sigma_3 dW_t^3, \end{aligned} \quad (3.2)$$

where x^2 expresses a state variable corresponding to the difference between medium-term and long-term commodity futures prices, and x^3 is a state variable corresponding to the long-term portion of the term structure. $W_t^i (i=1, 2, 3)$ have the following correlations of standard Brownian motions under an equivalent martingale measure (EMM):

$$dW_t^i \cdot dW_t^j = \rho_{ij} dt, \quad i, j = 1, 2, 3. \quad (3.3)$$

Parameter κ expresses x^1 's speed of reversion to $x^2 + x^3$, and γ expresses x^2 's speed of attenuation. If $\gamma > 0$, then x^2 is pulled back towards 0. β expresses the speed with which x^3 reverts to α/β when $\beta \neq 0$.

Therefore, intuitively, if $\kappa > \gamma > \beta > 0$, over the course of time the spot price x^1 (spot price) $\rightarrow x^2 + x^3$ (medium-term price) $\rightarrow x^3$ (long-term price) with the trend expressed.

The stochastic differential equations of individual state variables can be solved analytically and expressed as follows:

$$\begin{aligned}
 x_t^1 &= e^{-\kappa t} x_0^1 + \frac{\kappa}{\kappa - \gamma} (e^{-\gamma t} - e^{-\kappa t}) x_0^2 + \frac{\kappa}{\kappa - \beta} (e^{-\beta t} - e^{-\kappa t}) x_0^3 \\
 &\quad + \frac{\alpha}{\beta} \left(1 - \frac{\kappa}{\kappa - \beta} e^{-\beta t} + \frac{\beta}{\kappa - \beta} e^{-\kappa t} \right) + \sigma_1 \int_0^t e^{-\kappa(t-s)} dW_s^1 \\
 &\quad + \sigma_2 \frac{\kappa}{\kappa - \gamma} \int_0^t (e^{-\gamma(t-s)} - e^{-\kappa(t-s)}) dW_s^2 \\
 &\quad + \sigma_3 \frac{\kappa}{\kappa - \beta} \int_0^t (e^{-\beta(t-s)} - e^{-\kappa(t-s)}) dW_s^3, \\
 x_t^2 &= e^{-\gamma t} x_0^2 + \sigma_2 \int_0^t e^{-\gamma(t-s)} dW_s^2, \\
 x_t^3 &= e^{-\beta t} x_0^3 + \frac{\alpha}{\beta} (1 - e^{-\beta t}) + \sigma_3 \int_0^t e^{-\beta(t-s)} dW_s^3.
 \end{aligned} \tag{3.4}$$

From this the futures price is expressed as shown below.

Theorem 2.1: Using $G_T(t)$ to represent the price at time t of a futures with expiration T , under the EMM:

$$G_T(t) = E_t[S_T] = \exp \left\{ \mu_{11}(x_t^1, x_t^2, x_t^3, T-t) + \frac{\Sigma_{11}(T-t)}{2} \right\}. \tag{3.5}$$

In this equation, E_t denotes the conditional expectation at time t . For a discussion of μ_{11} and Σ_{11} see [Appendix A](#).

Proof: S_T has a log-normal distribution and the result can therefore be found by calculating the moment generating function of the normal distribution.

Next consider the market price of risk. $\theta(t) = (\theta_1(t), \theta_2(t), \theta_3(t))$ consists of the market price of risk for state variables x^1 , x^2 and x^3 . At this time, the following relationship holds true between the observed measure P and the equivalent Martingale measure Q :

$$W_t^P = W_t^Q + \int_0^t \theta(u) du. \tag{3.6}$$

Therefore, under measure P , the stochastic differential equations that are satisfied by individual state variables are

$$\begin{aligned}
 dx_t^1 &= \kappa(x_t^2 + x_t^3 - x_t^1)dt + \sigma_1 \theta_1(t)dt + \sigma_1 dW_t^{1,P} \\
 dx_t^2 &= -\gamma x_t^2 dt + \sigma_2 \theta_2(t)dt + \sigma_2 dW_t^{2,P}, \\
 dx_t^3 &= (\alpha - \beta x_t^3)dt + \sigma_3 \theta_3(t)dt + \sigma_3 dW_t^{3,P}.
 \end{aligned} \tag{3.7}$$

In particular, rewriting $\theta(t)$ with the state variables and time as $\theta(t, x^1, x^2, x^3)$:

$$\begin{aligned}
 \theta_1(t, x^1, x^2, x^3) &= -a(x_t^1 - x_t^2 - x_t^3), \\
 \theta_2(t, x^1, x^2, x^3) &= -bx_t^2, \\
 \theta_3(t, x^1, x^2, x^3) &= \begin{cases} c - dx_t^3, & (\beta \neq 0), \\ c & (\beta = 0). \end{cases}
 \end{aligned} \tag{3.8}$$

The stochastic differential equation (3.7) can therefore be rewritten as

$$\begin{aligned}
dx_t^1 &= \hat{\kappa}(x_t^2 + x_t^3 - x_t^1)dt + \sigma_1 dW_t^{1,P}, \\
dx_t^2 &= -\hat{\gamma}x_t^2 dt + \sigma_2 dW_t^{2,P}, \\
dx_t^3 &= (\hat{\alpha} - \hat{\beta}x_t^3)dt + \sigma_3 dW_t^{3,P},
\end{aligned} \tag{3.9}$$

where

$$\hat{\kappa} = \kappa + \sigma_1 a, \quad \hat{\gamma} = \gamma + \sigma_2 b, \quad \hat{\alpha} = \alpha + \sigma_3 c, \quad \hat{\beta} = \beta + \sigma_3 d.$$

Remark 2.1: In the discussion above, when $\beta = 0$, solving for the limit will give an analytic expression. Also, when $\beta = 0$, x^3 does not have a constant mean reversion level and the model itself does not have an ultimate mean reversion level. Below, we refer to cases where $\beta \neq 0$ as the ‘constant mean reversion model’, and $\beta = 0$ as the ‘stochastic mean reversion model’. Both types of models are essentially contained by Cortazar and Naranjo (2006) and Casassus and Collin-Dufresne (2005).

Remark 2.2: A two-factor constant mean reversion or a two-factor stochastic mean reversion model can be obtained by setting $x^2 \equiv 0$. These models are essentially the same as in Korn (2005), who used two-factor models for an analysis of hedging. For the two-factor model in the subsequent analysis, we place the restriction $x^2 \equiv 0$ on our three-factor models.

3.3 Estimation of Parameters

This section estimates the parameters in the model. Using v_n and w_n as white noise with mean 0 and variance 1, the model described above can be expressed as the following system and observation models.

System model:

$$\begin{aligned}
x_n &= F_n x_{n-1} + C_n^x + Q_n v_n, \\
x_n &= \begin{pmatrix} x^1 \\ x^2 \\ x^3 \end{pmatrix}, \\
F_n &= \begin{pmatrix} e^{-\hat{\kappa}\Delta t} & \frac{\hat{\kappa}}{\hat{\kappa} - \hat{\gamma}} \{e^{-\hat{\gamma}\Delta t} - e^{-\hat{\kappa}\Delta t}\} & \frac{\hat{\kappa}}{\hat{\kappa} - \hat{\beta}} \{e^{-\hat{\beta}\Delta t} - e^{-\hat{\kappa}\Delta t}\} \\ 0 & e^{-\hat{\gamma}\Delta t} & 0 \\ 0 & 0 & e^{-\hat{\beta}\Delta t} \end{pmatrix}, \\
E[Q_n] &= 0, \quad G_n = \text{Cov}[Q_n] = (\Sigma_{ij}(\Delta t)), \\
C_n^x &= \begin{pmatrix} \frac{\hat{\alpha}}{\hat{\beta}} \left\{ 1 - \frac{\hat{\kappa}e^{-\hat{\beta}\Delta t} - \hat{\beta}e^{-\hat{\kappa}\Delta t}}{\hat{\kappa} - \hat{\beta}} \right\} \\ 0 \\ \frac{\hat{\alpha}}{\hat{\beta}} \{1 - e^{-\hat{\beta}\Delta t}\} \end{pmatrix},
\end{aligned}$$

where $\Sigma_{ij}(\Delta t)$ is the covariance. For the specific formulae, see [Appendix A](#).

Observation model:

$$y_n = H_n x_n + C_n^y + R_n w_n,$$

$$y_n = \begin{pmatrix} \log G_{T_{1n}}(0) \\ \vdots \\ \log G_{T_{mn}}(0) \end{pmatrix}, H_n = \begin{pmatrix} e^{-\kappa T_{1n}} & \frac{\kappa}{\kappa - \gamma} \{e^{-\gamma T_{1n}} - e^{-\kappa T_{1n}}\} & \frac{\kappa}{\kappa - \beta} \{e^{-\beta T_{1n}} - e^{-\kappa T_{1n}}\} \\ \vdots & \vdots & \vdots \\ e^{-\kappa T_{mn}} & \frac{\kappa}{\kappa - \gamma} \{e^{-\gamma T_{mn}} - e^{-\kappa T_{mn}}\} & \frac{\kappa}{\kappa - \beta} \{e^{-\beta T_{mn}} - e^{-\kappa T_{mn}}\} \end{pmatrix},$$

$$R_n = \begin{pmatrix} h_1^2 & 0 \\ & \ddots \\ 0 & h_m^2 \end{pmatrix}, C_n^y = \begin{pmatrix} \frac{\Sigma_{11}(T_{1n})}{2} + \frac{\alpha}{\beta} \left\{ 1 - e^{-\kappa T_{1n}} - \frac{\kappa}{\kappa - \beta} \{e^{-\beta T_{1n}} - e^{-\kappa T_{1n}}\} \right\} \\ \vdots \\ \frac{\Sigma_{11}(T_{mn})}{2} + \frac{\alpha}{\beta} \left\{ 1 - e^{-\kappa T_{mn}} - \frac{\kappa}{\kappa - \beta} \{e^{-\beta T_{mn}} - e^{-\kappa T_{mn}}\} \right\} \end{pmatrix},$$

where R_n are the observational errors and $h_i, i = (1, \dots, m)$ are the standard deviations.

In light of the computational burden, we assume that the observational errors of futures at individual maturities are independent. The parameters are estimated using the Kalman filter of this state-space representation. More specifically, the following prediction and filtering are alternately repeated and a parameter set ϑ is obtained so as to maximize the log-likelihood.

Prediction:

$$x_{n|n-1} = F_n x_{n-1|n-1} + C_n^x,$$

$$V_{n|n-1} = F_n V_{n-1|n-1} F_n^t + G_n.$$

Filtering:

$$d_{n|n-1} = H_n V_{n|n-1} H_n^t + R_n,$$

$$K_n = V_{n|n-1} H_n^t d_{n|n-1}^{-1},$$

$$x_{n|n} = x_{n|n-1} + K_n (y_n - H_n x_{n|n-1} - C_n^y),$$

$$V_{n|n} = (I - K_n H_n) V_{n|n-1}.$$

Log-likelihood:

$$l(\vartheta) = -\frac{1}{2} \left\{ mN \log(2\pi) + \sum_{n=1}^N \log |\det(d_{n|n-1})| + \sum_{n=1}^N u_n^t d_{n|n-1}^{-1} u_n \right\},$$

$$u_n = y_n - H_n x_{n|n-1} - C_n^y.$$

Even if optimal values are not set for the initial values of x and V , as the calculation proceeds using the Kalman filter, both approach optimal values. Therefore, the initial value problem can be avoided by discarding several steps when estimating parameters without using the likelihood calculation. Estimations of two-factor models are obtained similarly.

3.3.1 Estimation Results

The constant mean reversion model and the stochastic mean reversion model parameters were estimated using the procedure described above. The following data were used for the estimations.

3.3.1.1 NYMEX WTI (light sweet crude oil)

Data in five business day increments were used for the periods January 1997–October 2002, January 1997–October 2003 and January 1997–November 2007; futures contracts were: Front Month, 1 Dec., 2 Dec., 3 Dec., 4 Dec., 5 Dec., 6 Dec. and 7 Dec. Here, j Dec. indicates the j -th month contract expiring in December. If the Front Month is 1 Dec., it was used as the front month. Data up to 10 Dec. exists after April 2007. However, data from 8 Dec. to 10 Dec. were not used in the estimation due to their lack of reliability.

3.3.1.2 LME Copper

Data in five business day increments were used for the periods September 2002–November 2004 and September 2002–December 2007; futures contract were: Front Month, 1 Dec., 2 Dec., 3 Dec., 4 Dec., 5 Dec. and 6 Dec. If the Front Month is 1 Dec., it was used as the front month.

These are liquid and typical assets of oil and metal futures. The choice of time period is the longest period for which the data have mid-term (7 Dec. in WTI, 6 Dec. in Copper) futures. [Tables 3.1](#) through [3.4](#) show the parameters and observational errors R_{it} obtained using the data described above.

Here, we note that observational errors in three-factor models are very small and that the model replicates the observed futures prices very well. For the two-factor constant mean reversion model, estimates of β using the WTI data up to 2003 and Copper data up to 2007 were $\beta \approx 0.4$, and hence, x^3 was observed to have a constant mean reversion level, but, in all other periods, both WTI and Copper had a β value of virtually 0. Therefore, x^3 does not fluctuate with a constant mean reversion level, but is rather more similar to a random walk. The parameters for Copper are significantly different between the data set up to 2004 and the data set up to 2007. When estimations are made using the data up to 2004, there is only a little more than 2 years of data used, and presumably the calculation results are biased parameters that are optimized to these 2 years. The market price of risk is expressed largely in parameters c and d for both WTI and Copper. We also observe that the standard errors of Copper's parameters are worse than those of WTIs, due, in part, to the lack of data used in the estimation. Finally, three-factor models show better fitting results than two-factor models in terms of AIC (Akaike's Information Criterion).

3.3.2 Comparison with Actual Data

This section verifies the degree of correlation between the state variables calculated with the Kalman filter using data through 2007 and settlement futures prices for NYMEX WTI and LME Copper.

As explained in [Section 3.2](#), the state variables correspond to the term structure of futures. In this case, the state variables are assumed to have the correspondences given in [Table 3.5](#) and the analysis seeks to determine the degree of correlation between them. [Table 3.6](#) shows the correlations for state variables and logarithmic prices calculated from five business day increments.

Both WTI and Copper generally have high correlations, indicating that the movement of state variables roughly corresponds to the actual data. Also, the three-factor models provide higher correlations than the two-factor models.

Next, we examine whether the models can reproduce the actual term structures of futures prices. [Figure 3.1](#) shows the term structures of two-factor and three-factor models against market prices of WTI futures on 3 November 2003, 1 November 2004, 1 November 2005, 1 November 2006 and 1 November 2007. [Figure 3.2](#) shows the results for Copper on 1 December 2003, 1 December 2004, 1 December 2005, 1 December 2006 and 3 December 2007. We can see in those cases that three-factor models can replicate the actual term structures well, whereas two-factor models have some difficulty in capturing the actual term structures. In particular, the difference in fitting between the two-factor model and the three-factor model frequently occurs in 2006 and 2007 to the extent that can be observed in [Figures 3.1](#) and [3.2](#).

TABLE 3.1 Three-Factor Model (WTI)

	Constant mean reversion model						Stochastic mean reversion model					
	Nov 07	Std Err	Oct 03	Std Err	Oct 02	Std Err	Nov 07	Std Err	Oct 03	Std Err	Oct 02	Std Err
κ	1.112	0.011	1.007	0.014	1.107	0.020	1.090	0.009	1.159	0.015	1.048	0.017
γ	0.275	0.007	0.159	0.012	0.293	0.018	0.262	0.007	0.284	0.009	0.253	0.011
α	0.006	0.010	0.052	0.023	0.311	0.024	-0.009	0.001	-0.007	0.001	-0.007	0.002
β	0.004	0.001	0.021	0.008	0.110	0.007	—	—	—	—	—	—
a	0.000	—	0.000	—	0.597	1.787	0.000	—	0.000	—	0.529	1.861
b	0.000	—	0.000	—	0.000	—	0.000	—	0.000	—	0.000	—
c	0.528	0.341	0.000	—	0.070	0.219	0.482	0.266	0.062	0.318	0.063	0.347
d	0.005	0.019	0.002	0.039	0.017	0.071	—	—	—	—	—	—
Sigma_1	0.362	0.011	0.359	0.014	0.371	0.015	0.3630	0.010	0.387	0.015	0.363	0.013
Sigma_2	0.144	0.004	0.202	0.028	0.374	0.043	0.142	0.003	0.137	0.004	0.138	0.005
Sigma_3	0.170	0.059	0.251	0.067	0.399	0.065	0.190	0.005	0.171	0.006	0.182	0.008
$\text{Rho}_{[12]}$	0.155	0.047	0.048	0.077	-0.173	0.051	0.153	0.047	0.155	0.065	0.192	0.070
$\text{Rho}_{[23]}$	-0.601	0.205	-0.851	0.199	-0.983	0.116	-0.539	0.033	-0.599	0.052	-0.619	0.062
$\text{Rho}_{[31]}$	0.396	0.141	0.237	0.067	0.352	0.065	0.369	0.034	0.282	0.047	0.261	0.051
Front Month	0.041	0.001	0.049	0.002	0.047	0.002	0.041	0.001	0.047	0.002	0.045	0.002
1Dec	0.000	—	0.004	0.002	0.000	—	0.000	—	0.000	—	0.000	—
2Dec	0.007	0.000	0.013	0.001	0.006	0.000	0.007	0.000	0.008	0.000	0.007	0.000
3Dec	0.002	0.000	0.008	0.000	0.003	0.000	0.002	0.000	0.002	0.000	0.002	0.000
4Dec	0.004	0.000	0.001	0.000	0.004	0.000	0.004	0.000	0.005	0.000	0.005	0.000
5Dec	0.003	0.000	0.001	0.000	0.002	0.000	0.003	0.000	0.003	0.000	0.003	0.000
6Dec	0.000	—	0.001	0.000	0.002	0.000	0.000	—	0.000	—	0.000	—
7Dec	0.004	0.000	0.003	0.000	0.004	0.000	0.004	0.000	0.004	0.000	0.004	0.000
AIC	-27157.30		-16837.37		-14499.01		-27146.66		-16701.48		-14241.34	

TABLE 3.2 Three-Factor Model (Copper)

	Constant mean reversion model				Stochastic mean reversion model			
	Dec 07	Std Err	Nov 04	Std Err	Dec 07	Std Err	Nov 04	Std Err
κ	0.766	0.046	0.918	0.173	0.740	0.039	0.930	0.180
γ	0.177	0.036	0.143	0.069	0.161	0.024	0.153	0.036
α	0.329	0.114	0.036	0.633	-0.059	0.010	-0.014	0.009
β	0.052	0.015	0.007	0.084	—	—	—	—
a	0.000	—	0.000	—	0.000	—	0.000	—
b	0.000	—	1.720	2.915	0.000	—	2.044	1.890
c	0.092	0.212	0.079	0.567	0.153	0.331	0.096	0.546
d	0.000	—	0.000	0.023	—	—	—	—
Sigma_1	0.268	0.013	0.224	0.021	0.268	0.013	0.224	0.021
Sigma_2	0.631	0.199	0.317	0.307	0.437	0.043	0.287	0.056
Sigma_3	0.658	0.201	0.319	0.350	0.440	0.039	0.287	0.046
$\text{Rho}_{\{12\}}$	0.074	0.080	0.368	0.339	0.256	0.069	0.402	0.112
$\text{Rho}_{\{23\}}$	-0.929	0.057	-0.840	0.593	-0.845	0.033	-0.801	0.072
$\text{Rho}_{\{31\}}$	0.237	0.082	0.130	0.250	0.200	0.070	0.155	0.139
Front Month	0.000	—	0.000	—	0.000	—	0.000	—
1Dec	0.007	0.000	0.006	0.001	0.007	0.000	0.006	0.001
2Dec	0.003	0.000	0.001	0.001	0.003	0.000	0.001	0.001
3Dec	0.005	0.000	0.004	0.000	0.005	0.000	0.003	0.000
4Dec	0.000	—	0.003	0.001	0.000	—	0.003	0.001
5Dec	0.010	0.001	0.002	0.002	0.010	0.001	0.002	0.002
6Dec	0.011	0.000	0.008	0.002	0.013	0.001	0.008	0.001
AIC	-9538.94		-4259.00		-9520.27		-4264.08	

TABLE 3.3 Two-Factor Model (WTI)

	Constant mean reversion model						Stochastic mean reversion model					
	Nov 07	Std Err	Oct 03	Std Err	Oct 02	Std Err	Nov 07	Std Err	Oct 03	Std Err	Oct 02	Std Err
κ	0.218	0.004	0.042	0.003	0.511	0.007	0.334	0.006	0.379	0.008	0.395	0.007
α	-0.074	0.017	1.187	0.033	0.234	0.014	-0.009	0.001	-0.003	0.001	-0.005	0.001
β	0.012	0.001	0.390	0.010	0.084	0.002	—	—	—	—	—	—
a	0.000	—	0.000	—	0.000	—	0.000	—	0.000	—	0.000	—
c	0.072	0.434	2.093	1.300	0.080	0.862	0.218	0.293	0.035	0.389	0.000	—
d	0.001	0.002	1.617	0.698	0.000	0.262	—	—	—	—	—	—
Sigma_1	0.426	0.027	0.171	0.009	0.225	0.012	0.193	0.006	0.175	0.008	0.190	0.010
Sigma_3	0.520	0.036	0.865	0.129	0.222	0.076	0.180	0.005	0.160	0.005	0.167	0.006
$\text{Rho}_{\{31\}}$	0.941	0.034	-0.748	0.145	0.114	0.061	0.540	0.029	0.419	0.041	0.375	0.047
Front Month	0.145	0.006	0.137	0.007	0.009	0.005	0.145	0.006	0.142	0.008	0.127	0.007
1Dec	0.101	0.009	0.087	0.009	0.002	0.003	0.075	0.005	0.075	0.006	0.066	0.005
2Dec	0.029	0.002	0.016	0.001	0.000	0.001	0.016	0.001	0.016	0.001	0.015	0.001
3Dec	0.009	0.000	0.002	0.000	0.000	0.000	0.002	0.000	0.003	0.000	0.002	0.001
4Dec	0.001	0.000	0.003	0.000	0.000	0.000	0.003	0.000	0.005	0.000	0.006	0.000
5Dec	0.002	0.000	0.000	—	0.000	—	0.000	—	0.004	0.000	0.004	0.000
6Dec	0.000	—	0.004	0.000	0.000	0.000	0.005	0.000	0.000	—	0.000	—
7Dec	0.004	0.000	0.007	0.000	0.000	0.001	0.009	0.000	0.005	0.000	0.005	0.000
AIC	-22211.90		-14422.80		-12770.44		-22632.97		-14387.29		-12253.04	

TABLE 3.4 Two-Factor Model (Copper)

	Constant mean reversion model				Stochastic mean reversion model			
	Dec 07	Std Err	Nov 04	Std Err	Dec 07	Std Err	Nov 04	Std Err
κ	0.122	0.006	0.327	0.016	0.137	0.018	0.327	0.015
α	2.960	0.151	0.002	0.124	−0.032	0.014	−0.016	0.003
β	0.407	0.021	0.001	0.017	—	—	—	—
a	0.000	—	0.000	—	0.000	—	0.000	—
c	1.879	5.062	0.000	—	0.484	0.497	0.000	—
d	0.152	0.652	0.034	0.708	—	—	—	—
Sigma1	0.275	0.015	0.223	0.021	0.265	0.013	0.218	0.020
Sigma3	0.927	0.469	0.092	0.212	0.469	0.049	0.194	0.015
Rho{31}	−0.239	0.138	0.893	2.021	0.014	0.078	0.409	0.079
Front Month	0.026	0.001	0.014	0.003	0.040	0.002	0.014	0.003
1Dec	0.000	—	0.003	0.003	0.003	0.004	0.003	0.003
2Dec	0.010	0.000	0.006	0.002	0.026	0.003	0.006	0.002
3Dec	0.000	—	0.005	0.001	0.018	0.003	0.005	0.001
4Dec	0.015	0.002	0.003	0.001	0.000	—	0.003	0.001
5Dec	0.031	0.004	0.005	0.001	0.024	0.002	0.005	0.001
6Dec	0.042	0.004	0.008	0.002	0.061	0.006	0.008	0.002
AIC	−7662.46		−3927.20		−6806.08		−3922.25	

TABLE 3.5 Correspondence of State Variables

WTI			Copper
3factor	x^1	Front Month futures Price	Front Month futures Price
	x^2	(3 DEC futures Price) − (6 DEC futures Price)*	(2 DEC futures Price) − (5 DEC futures Price)*
	x^3	6 DEC futures Price	5 DEC futures Price
2factor	x^1	2 DEC futures Price	Front Month futures Price
	x^3	6 DEC futures Price	5 DEC futures Price

* x^2 is compared with the spread between 6 Dec and 3 Dec for WTI, and the spread between 5 Dec and 2 Dec for Copper.

TABLE 3.6 Correlations

		Constant mean reversion model		Stochastic mean reversion model	
		WTI	Copper	WTI	Copper
3factor	x^1	0.926	0.923	0.927	0.924
	x^2	0.944	0.933	0.941	0.939
	x^3	0.980	0.841	0.977	0.821
2factor	x^1	0.857	0.885	0.864	0.876
	x^3	0.959	0.720	0.986	0.784

3.4 Futures Hedging Techniques

This section describes a method for constructing a hedging strategy for one unit of a long-term futures contract and describes how the three-factor model used in this paper can be applied to this task.

The equation expressing the futures price uses state variables x^1 , x^2 and x^3 so that the shape of the futures price changes according to changes in these state variables (assuming no change in the parameters). Therefore, it is possible in theory to hedge against long-term futures price fluctuations by calculating the deltas of the state variables for the long-term futures price and taking a position $\Phi = (\phi_1, \phi_2, \phi_3)^t$ in the nearer maturity futures that cancels out those deltas.




FIGURE 3.1 WTI futures term structure.

In a three-factor model, there are three factors to be hedged and therefore, futures with three different expiration maturities will be required to build the hedge portfolio. $G_{T_1}(t)$, $G_{T_2}(t)$ and $G_{T_3}(t)$ express nearer maturity futures prices of different expirations, and $G_{T_4}(t)$ the long-term futures price to be hedged. In this case, Φ is the solution to the following simultaneous equation:

$$A\Phi = b,$$

where

$$A = \begin{pmatrix} \frac{\partial G_{T_1}(t)}{\partial x_1} & \frac{\partial G_{T_2}(t)}{\partial x_1} & \frac{\partial G_{T_3}(t)}{\partial x_1} \\ \frac{\partial G_{T_1}(t)}{\partial x_2} & \frac{\partial G_{T_2}(t)}{\partial x_2} & \frac{\partial G_{T_3}(t)}{\partial x_2} \\ \frac{\partial G_{T_1}(t)}{\partial x_3} & \frac{\partial G_{T_2}(t)}{\partial x_3} & \frac{\partial G_{T_3}(t)}{\partial x_3} \end{pmatrix}, \quad b = \begin{pmatrix} \frac{\partial G_{T_4}(t)}{\partial x_1} \\ \frac{\partial G_{T_4}(t)}{\partial x_2} \\ \frac{\partial G_{T_4}(t)}{\partial x_3} \end{pmatrix}.$$

This paper refers to hedging using the hedging portfolio Φ as a 'delta hedge'. For a two-factor model, it is possible to construct a delta hedge in a similar way by eliminating the second factor x^2 in the corresponding three-factor model. We verify the degree of hedging error against this hedging portfolio when time series data are applied. For the purposes of this paper, the 'hedging error rate' is expressed as the final cumulative hedging error divided by the price of the instrument to be hedged at the time the hedge commences.




FIGURE 3.2 Copper futures term structure.

For comparison, we calculate the hedging error ratio for hedges such as those performed by Metallgesellschaft in which an equivalent number of nearer maturity futures is held against the futures to be hedged. This paper refers to this hedging method as the ‘parallel hedge’. Metallgesellschaft hedged its long-term futures with extremely short-term futures of one to three contract months. However, given the increased liquidity of current commodities futures markets into the medium-term range, we verify the effectiveness of parallel hedges using futures of up to 6 years for WTI and up to 5 years for Copper.

Unless specifically stated to the contrary, the discussion below refers to hedges against 10 Dec. from the front month for the WTI and 8 Dec. for Copper, for which prices are estimated by our models. For the hedging period, it is assumed that the position will be closed with an offsetting trade of a 6 Dec. futures for the WTI. In other words, a 4-year hedge is entered into that reduces the time to maturity of the instrument to be hedged from 10 to 6 years. For Copper, it is assumed that the position is closed with an offsetting trade of a 5 Dec. futures, resulting in a 3-year hedge that reduces the time to maturity of the instrument to be hedged from 8 to 5 years. For the parallel hedge, futures for the listed Decs. are used as hedge assets. For the delta hedges of three-factor models, 1, 4 and 6 Dec., and 1, 3 and 5 Dec. are used for WTI and Copper, respectively. For the delta hedges of two-factor models, 4 and 6 Dec. and 3 and 5 Dec. are used for WTI and Copper, respectively. Positions in each futures contract month are adjusted on the first business day of the month after reviewing hedging ratios each month. For both the parallel hedge and delta hedge, upon the elapse of 1 year, positions are rolled to the same contract month in the next year. (For example, if a Dec. 6 position is used to initiate a hedge on Dec. 12, after the elapse of 1 year, the Dec. 6 position used in the hedge will be rolled over to Dec. 7.) Liquidity declines the more distant the futures, but Dec. futures have comparatively high liquidity, and given the infrequency with which hedge

ratios are changed and the small degree of change in the number of units required for hedging, this is considered a realistic hedge. In selecting futures contract months, this analysis uses combinations that provide the relatively small hedging error rates obtained in [Section 3.6](#). A similar procedure is used to select futures contract months for two-factor models.

3.4.1 Hedging Error Rate of The Parallel Hedge

We first verify the hedging error rate achieved using the parallel hedge. The price of the futures contract month to be hedged is calculated based on the constant mean reversion model and data up to 2007 using parameters and state variables estimated with the Kalman filter.

Figure 3.3 shows the cumulative hedging error rates (the cumulative hedging error divided by the price of the instrument to be hedged at the time the hedge commences) using WTI and Copper time series data. In this case, for the Front Month the cumulative hedging error rate is expressed for a parallel hedge rolled over to the next-expiring contract month each month; for others, the cumulative hedging error rate is expressed for a parallel hedge with a one-year roll using Decs. for each year. The futures to be hedged are the WTI Dec. 13 and the Copper Dec. 12 and the hedge terminates at the most recently available data (2007). The horizontal axis expresses the amount of time elapsed since the commencement of the hedge; the vertical axis, the cumulative hedging error rate. The same notation is used for other graphs in this paper.

As can be observed from Figure 3.3, the error is smaller the more distant the futures used to hedge. Copper has a larger hedging error rate than WTI, indicating that the components in Copper's term structure that change in parallel are smaller than WTIs. However, even using the most distant futures with the smallest hedging error rate, the hedging error rates with a parallel hedge were still approximately 12% for WTI and approximately 32% for Copper.

3.4.2 Hedging Error Rate of The Delta Hedge

This section determines the hedging error rate for delta hedges for both the constant mean reversion model and the stochastic mean reversion model.

For each model, parameters estimated from data up to 2007 were used, and, for verification purposes, two methods were used to estimate the state variables in order to estimate long-term futures prices. In the first method, state variables were estimated using the Kalman filter ('Kalman filter state variables' hereafter); the second estimation created simultaneous linear equations for the state variables so that the futures price of the model matches the futures price of the futures contract month to be hedged, allowing state variables to be calculated by solving these equations ('simultaneous equation-based state variables' hereafter). To compare the relative precision of hedging using the model described in this paper, the verifications below give the results for the most distant futures, which was the most precise for the parallel hedge. Notation follows the practice used for parallel hedges.




FIGURE 3.3 Cumulative hedging error rates of the parallel hedge.

Verifications were performed with different futures to be hedged, and the observed hedging error rates are summarized in Table 3.7. The parallel hedge is able to provide effective hedging when the overall term structure changes in parallel, but generates large hedging errors when there are changes in the shape of the term structure. By contrast, the delta hedge works much better than the parallel hedge (see Table 3.7). Two- and three-factor models provide relatively similar results, although three-factor models work better for Copper. However, two-factor models have some difficulty in replicating actual term structures, as shown in Figures 3.1 and 3.2.

Comparing Kalman filter state variables and simultaneous equation-based state variables when performing a delta hedge, for Copper the estimation of Kalman filter state variables produces a large hedging error during the term of the hedge, as can be seen in Figures 3.4 and 3.5. This is presumably due to model prices not being obtained in a manner consistent with both the asset prices used in the hedge and the prices of the assets to be hedged. When using simultaneous equation-based state variables, model prices (excluding rollover timing) match the prices of the assets used in the hedge and the assets to be hedged. On the other hand, when using Kalman filter state variables, the actual prices of the assets used in the hedge differ from the model prices, resulting in a hedging error when hedging is performed. For WTI, the observational error was small for the futures contract month used in the hedge and virtually equivalent to the simultaneous equation-based state variables, indicating that there is little difference due to the method by which state variables are determined.

Because of the above result, the discussion below uses only state variables that are calculated by solving simultaneous equations for both WTI and Copper models.

TABLE 3.7 Hedging Error Rates

		Constant mean reversion model					Stochastic mean reversion model				
		Kalman Filter		Equations Based		Parallel	Kalman Filter		Equations Based		Parallel
		3factor	2factor	3factor	2factor		3factor	2factor	3factor	2factor	
		(%)	(%)	(%)	(%)		(%)	(%)	(%)	(%)	
WTI	DEC 07	1.0	1.3	0.7	1.3	4.0	0.7	0.8	0.4	0.4	3.8
	DEC 08	1.0	0.6	0.9	0.6	8.4	0.7	1.5	0.5	0.7	8.3
	DEC 09	0.8	0.8	1.0	0.8	3.4	0.5	-0.8	0.8	0.6	3.2
	DEC 10	-1.5	-2.3	-2.5	-2.5	-3.5	-1.7	-0.5	-2.7	-2.5	-4.0
	DEC 11	-3.1	-3.8	-3.2	-3.8	2.7	-3.1	-2.8	-3.1	-2.6	2.4
	DEC 12	-3.2	-4.2	-3.7	-4.2	8.0	-3.0	0.4	-3.6	-0.9	7.8
Copper	DEC 13	-0.8	-1.1	0.0	-1.1	12.1	-0.3	0.5	0.5	3.2	12.3
	DEC 10	-0.4	-1.2	-0.5	-2.0	12.4	-0.1	0.0	0.0	-1.8	12.4
	DEC 11	0.0	-5.2	0.2	-4.8	30.3	1.2	2.5	1.4	1.4	30.3
	DEC 12	-3.1	-20.8	-6.7	-20.8	32.1	0.9	4.8	-3.2	3.7	32.0




FIGURE 3.4 Cumulative hedging error rates of the delta hedge (constant mean reversion model).




FIGURE 3.5 Cumulative hedging error rates of the delta hedge (stochastic mean reversion model).

3.5 Stability of The Delta Hedge

The verifications so far have estimated the parameters based on data that included the entire hedge period. However, in actual practice, the parameter estimation period and the hedge period differ. Discussions so far have also assumed that the hedges target long-term futures, but general practice is for long-term contracts to be forwards rather than futures, which requires that interest-rate factors also be taken into account.

In this section, we confirm the following settings so as to conduct the verifications in a state as close as possible to actual practice.

1. No overlap between the parameter estimation period and the hedge period.
2. Hedging against forwards.

In this paper, cases in which the entire hedge period is included in the parameter estimation period are referred to as 'In Sample', while hedges in which there are separate parameter estimation periods are referred to as 'Out of Sample'.

3.5.1 Out of Sample Hedges

To verify the effectiveness of the Out of Sample hedge, this section uses the parameters estimated in Section 3.3 with the data through 2002 or 2003 for WTI and through 2004 for Copper. The futures to be hedged is Dec. 12 (hedge period from 2002 to 2006) or Dec. 13 (hedge period from 2003 to 2007) for WTI and Dec. 12 (hedge period from 2004 to 2007) for Copper.

Table 3.8 summarizes the results of the verifications using time series data for hedging error rates when hedging long-term futures prices as estimated with the model using these parameters. For purposes of comparison, we have also listed the results for the In Sample estimations in the previous section.

For comparison between the two-factor models and the three-factor models, it is observed that the model producing the smaller absolute error In Sample also creates the smaller absolute error Out of Sample for all cases of Table 3.8.

For the three-factor stochastic mean reversion model, we confirmed that the differences in hedging error rates due to differences in the parameter estimation period were not that large for WTI. For Copper, there were differences in In Sample and Out of Sample hedging error rates, in part due to the differences in the parameters obtained for In Sample and Out of Sample. However, the error is not extreme and even though there are differences in the parameters obtained using the maximum likelihood method, hedging under the model is considered to be relatively stable. On the other hand, for the two-factor stochastic mean reversion model, there are more differences between In sample and Out of sample hedging error rates: this tendency seems stronger for Copper than for WTI.

For the constant mean reversion models, in the Out of Sample WTI, long-term price levels changed due to the sharp increases in oil prices beginning in 2003, virtually eliminating mean reversion. Nonetheless, the long-term futures prices calculated by the models revert to the mean levels observed

TABLE 3.8 Hedging Error Rates of the Delta Hedge (Out of Sample)

		Constant mean reversion model		Stochastic mean reversion model	
		In sample (%)	Out of sample (%)	In sample (%)	Out of sample (%)
WTI	3factor (DEC 12)	-3.7	-32.1	-3.6	-3.4
	2factor (DEC 12)	-4.2	-33.5	-0.9	1.5
	3factor (DEC 13)	0.0	-7.6	0.5	0.2
	2factor (DEC 13)	-1.1	-16.5	3.2	5.3
Copper	3factor (DEC 12)	-6.7	-6.9	-3.2	-1.9
	2factor (DEC 12)	-20.8	10.0	3.7	11.0

in the data through 2002 or 2003, increasing the hedging error rates. In light of this, it is likely that the constant mean reversion model is more prone to hedging error when there are changes in mean reversion levels, etc., indicating that it is better to use the stochastic mean reversion model when prices are based on the hedge.

Due to the above results, the discussion in the next subsection uses only the three-factor stochastic mean reversion model.

3.5.2 Hedging Long-Term Forward Contracts

The discussions to this point have assumed that futures would be hedged, but common practice is to trade forwards for the long-term portion that is not traded on exchanges. If interest rates are deterministic or move independently from underlying assets, prices are the same for futures and forwards, but if interest rates are not deterministic, hedges must take account of their movements.

For simplicity, this discussion assumes that interest rates and underlying assets are independent, and describes hedging techniques when the instrument to be hedged is a forward and the assets used in the hedge are futures. The utility of this hedging technique is then verified using time series data.

We consider hedging long-term forwards with short-term futures in the following two steps.

1. Long-term forwards are hedged using long-term futures with the same expiration.
2. Long-term futures are hedged using the delta hedge with the nearer maturity futures as described in [Section 3.4](#).

Because Step 2 is explained in [Section 3.4](#), we explain the hedging technique of Step 1. The notation used is as follows: $F_T(t)$ is the price at point in time t of the forward with expiration T , $G_T(t)$ is the price at point in time t of the futures with expiration T and $P_T(t)$ is the price at point in time t of zero-coupon bond with expiration T .

In addition, $0 = t_0 < t_1 < \dots < t_m = T$. The amount of change in the forward profit/loss at point in time T during the period from point in time t_i through t_{i+1} is as follows: PV at point in time t_{i+1} is expressed as $(F_T(t_{i+1}) - F_T(0))P_T(t_{i+1})$. Therefore, the amount of change in PV for the forward during the period from t_i through t_{i+1} is

$$(F_T(t_{i+1}) - F_T(0))P_T(t_{i+1}) - (F_T(t_i) - F_T(0))P_T(t_i). \quad (3.10)$$

In this case, (3.10) can be reformed as follows:

$$\begin{aligned} (F_T(t_{i+1}) - F_T(0))P_T(t_{i+1}) - (F_T(t_i) - F_T(0))P_T(t_i) &= (F_T(t_{i+1}) - F_T(t_i))P_T(t_i) \\ &\quad + (F_T(t_i) - F_T(0))(P_T(t_{i+1}) - P_T(t_i)) \\ &\quad + (F_T(t_{i+1}) - F_T(t_i))(P_T(t_{i+1}) - P_T(t_i)). \end{aligned} \quad (3.11)$$

Thus ΔC_m , which denotes the accumulation of the last term on the right-hand evaluated at time T side of (3.11), is given by

$$\Delta C_m \approx \sum_{j=0}^{m-1} e^{\sum_{i=j+1}^{m-1} r_i(t_{i+1}-t_i)} (F_T(t_{j+1}) - F_T(t_j))(P_T(t_{j+1}) - P_T(t_j)), \quad (3.12)$$

where the instantaneous interest rate in each period $[t_i, t_{i+1}]$ is approximated as a constant r_i . Below, we ignore ΔC_m because it expresses a negligible amount corresponding to the quadratic variation.

Assuming that the interest rate is independent of the underlying asset prices, the forward price and futures price are equivalent ($F_T(t) = G_T(t)$) and Equation (3.11) can be expressed as

$$\begin{aligned} (F_T(t_{i+1}) - F_T(t_i))P_T(t_i) + (F_T(t_i) - F_T(0))(P_T(t_{i+1}) - P_T(t_i)) &= (G_T(t_{i+1}) - G_T(t_i))P_T(t_i) \\ &+ (G_T(t_i) - G_T(0))(P_T(t_{i+1}) - P_T(t_i)). \end{aligned} \quad (3.12)$$

The first term on the right-hand side expresses the change in the future, according to which a delta hedge is made using futures for the three nearer maturity contract months under the method described in Section 3.4. As a result, one unit of forwards can be given a proximate hedge using a portfolio comprising futures and zero-coupon bonds, as shown below.

1. A delta hedge using futures for the three nearer maturity contract months to hedge $P_T(t)$ units of futures $G_T(t)$.
2. Purchase of $G_T(t) - G_T(0)$ units of zero-coupon bond $P_T(t)$.

This hedging strategy is essentially the same as Schwartz's hedging strategy (see Schwartz 1997, pp. 963–964).

We analysed the hedging error rate when a 10-year WTI forward contract is hedged according to the above method using WTI futures and zero-coupon bonds for 4 years. Here, forward prices are assumed to be equal to theoretical futures prices calculated by our model. Note that the funds for the purchase of zero-coupon bonds and the cash flow generated by marking futures to market are invested/raised in short-term interest rates. For verification, we used In Sample parameters and, for simplification, the calculations of the zero-coupon bonds used the 8-year swap rate as spot yield; the calculations of short-term interest rates for investments and funding used the 1M LIBOR, and assume that the interest rate is independent of asset prices, thus forward price is equal to futures price. Table 3.9 shows the hedging error rates due to differences in the forwards to be hedged.

It can be seen that, in this verification, which used time series data for interest rates and futures prices, even assuming interest rates and underlying assets to be independent, the use of zero-coupon bonds and futures to hedge forwards was able to hedge virtually all of the interest rate factors generated by the difference between forwards and futures. However, if interest rates are not hedged, there are cases where large hedging errors are generated during the hedge period, as can be seen from Figure 3.6, so the idea that there does not need to be a hedge on the interest-rate portion is not supported.

TABLE 3.9 Hedging Error Rates of the Delta Hedge

	DEC07	DEC08	DEC09	DEC 10	DEC 11	DEC 12	DEC 13
Using futures and bonds	0.2%	0.3%	0.4%	−2.3%	−2.6%	−3.1%	−0.2%
Using futures	0.1%	−1.3%	−2.4%	−9.2%	−5.8%	−10.0%	−5.6%




FIGURE 3.6 Cumulative hedging error rates of the delta hedge.

3.6 Measuring The Distribution of Hedging Error Rates

The analysis of hedging error rates based on time series data are limited to one of many possible paths. Thus, this section provides a simulation analysis to measure the distribution of hedging error rates resulting from fluctuations in the underlying assets. Three-factor stochastic mean reversion models are used in the simulation, where the In Sample parameters are used, and futures are hedged by futures with shorter maturities. We list below the specific procedures for the simulation.

1. Historical daily futures prices were created based on parameters and state variables estimated using the Kalman filter.
2. The error rate between the futures prices based on the model and created in Step 1 vs. the actual futures prices quoted on exchanges (for WTI, the front month and 1–6 Dec.; for Copper, the front month and 1–5 Dec.) was calculated $[(\text{Actual data} - \text{Model price}) / \text{Model price}]$ and then the mean and covariance of the error rate were obtained. Here, it was assumed that the error follows a multidimensional normal distribution.
3. Three-dimensional normal random numbers were created and the state variables were made to fluctuate according to the model so as to create a term structure for futures.
4. Multidimensional normal random numbers according to the distribution described in Step 2 were created for the futures term structure developed in Step 3, multiplied as error and added to the original term structure (Step 3 model prices + Step 3 model prices $\times$ Random numbers following the Step 2 error distribution).
5. Simultaneous equation-based state variables were calculated on the assumption that the term structure created in Step 4 was the term structure actually observed in the market.
6. The term structure created in Step 5 was used to estimate long-term prices, hedges were taken against those prices, and the final hedging error rate measured.
7. Steps 3 to 6 were repeated for a constant number of times to find the sample mean and sample standard deviation of the hedging error rates obtained.

3.6.1 Distribution of Hedging Error Rates

We performed 5000 trials for each combination of futures contract used in the hedge according to the procedures outlined above and calculated the average and standard deviation of the hedging error rates. There was little difference in the hedging error rates due to differences in initial values, so as hedged assets we used Dec. 17 for WTI for a period of 4 years beginning 2007 and Dec. 15 for Copper for a period of 3 years beginning 2007.

TABLE 3.10 Averages and Standard Deviations of Hedging Errors

Contract Months	Average (%)	Standard Deviation (%)	Contract Months	Average (%)	Standard Deviation (%)
WTI					
1-2-3	-22.6	20.9	2-3-5	-0.3	4.5
1-2-4	-10.0	10.3	2-3-6	0.3	2.2
1-2-5	-4.3	5.0	2-4-5	-0.1	4.9
1-2-6	-1.7	2.7	2-4-6	0.4	2.2
1-3-4	-4.9	8.0	2-5-6	0.4	1.9
1-3-5	-1.5	3.9	3-4-5	-0.3	6.0
1-3-6	-0.3	2.0	3-4-6	0.3	2.5
1-4-5	-0.5	4.3	3-5-6	0.3	2.0
1-4-6	0.1	2.0	4-5-6	0.3	2.5
1-5-6	0.3	1.8	Parallel Hedge		
2-3-4	-3.0	9.3	6Y Parallel	-0.5	6.0
Copper					
1-2-3	-14.7	11.5	2-3-4	0.3	8.5
1-2-4	-4.0	5.9	2-3-5	0.6	4.1
1-2-5	-1.1	3.7	2-4-5	-0.2	3.9
1-3-4	-0.7	6.8	3-4-5	-1.0	4.7
1-3-5	0.1	3.7	Parallel Hedge		
1-4-5	-0.2	3.7	5Y Parallel	0.4	16.3

See [Appendix B](#) for more on the number of futures units required due to differences in the selection of futures months at the time the hedge is commenced.

Table 3.10 shows the means and standard deviations for the obtained hedging error rates. The ‘contract months’ column refers to which Dec. from the front month is used for the hedge. For example, 1–4–6 refers to the hedge using the 1, 4 and 6 Dec. values. Likewise, ‘6Y Parallel (5Y Parallel)’ expresses the hedging error rate when a parallel hedge is entered into using the 6 (5) Dec. The price of the hedged assets for the parallel hedge is in the price found using 1–4–6 for WTI and 1–3–5 for Copper.

According to the results of the simulation, the appropriate selection of the futures contract months for the hedge portfolio when entering into a delta hedge has the potential for a more accurate hedge than the use of a parallel hedge. In particular, in terms of the required amounts (see [Appendix B](#)) and also the relationship between the means and standard deviations of hedging error rates, hedges for the WTI and Copper exhibited efficiency using 1, 4 and 6 Dec. and 1, 3 and 5 Dec., respectively. As a more general result, it was found that the futures used to create a hedge portfolio should, to the extent possible, have mutually disparate contract months. This is because when the futures used in the hedge are close to each other the x^1 , x^2 and x^3 delta structures are similar and a greater number of futures units is required to offset the delta of the instrument hedged. The greater the number of futures units used in the hedge, the larger the error expressed as differences in the prices of the model and actual futures. Conversely, when the futures are further away from each other, they have disparate delta structures, making it more likely that hedging will not require as many units.

3.7 Conclusion

This paper has demonstrated that using a three-factor Gaussian model with appropriate estimation of parameters, it is possible to reproduce the term structures of listed commodities futures (NYMEX WTI and LME Copper) during the time period studied and that long-term futures prices can be obtained that are consistent with liquid nearer maturity contracts. It was also found that two-factor Gaussian models have some difficulty in capturing the actual term structures of futures.

Furthermore, we have proposed a hedging technique for long-term futures and forwards contracts, comparing the results from this technique with the results from the simple short-term futures-based hedging strategy used by Metallgesellschaft (parallel hedge), and have verified that our proposed strategy is stable in many different circumstances (backwardation, contango, rising prices, declining prices, etc.).

In addition, it was found that a stochastic mean reversion model offered more stable hedging than a model with a constant mean reversion level. Also, it was observed that the model producing the smaller absolute error In Sample also creates the smaller absolute error Out of Sample.

We then used simulation to measure the hedging error rates obtained due to differences in the contract months of the futures used in the hedge. It was found that the futures used to create a hedge portfolio should, to the extent possible, have mutually disparate contract months.

In summary, the three-factor model with stochastic mean reversion is useful in practice for pricing long-term futures/forward contracts and for hedging them with appropriate selected liquid instruments.

Future issues include the evaluation of option values using the model and structuring of relevant hedging techniques. Commodities generally have average-based options which makes calculation complex. It would be useful to verify the hedging techniques and their efficiency.

Acknowledgements

We thank the two anonymous referees for their efforts and comments. The views expressed in this paper are those of the authors and do not necessarily represent the views of Mizuho-DL Financial Technology Co., Ltd.

References

- Bessembinder, H., Coughenour, J., Sequin, P. and Smoller, M., Mean reversion in equilibrium asset prices: evidence from the future term structure. *J. Finance*, 1995, **50**, 361–375.
- Black, F., The pricing of commodity contracts. *J. Financial Econ.*, 1976, **3**, 167–179.
- Brenann, M.J. and Crew, N.I., Hedging long-maturity commodity commitments with short-dated futures contracts. In *Mathematics of Derivatives Securities*, edited by M. Dempster and S. Pliska, pp. 165–190, 1997 (Cambridge University Press: Cambridge).
- Bühler, W., Korn, O. and Schöbel, R., Hedging long-term forwards with short-term futures: a two-regime approach. *Rev. Deriv. Res.*, 2004, **7**, 185–212.
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Cortazar, G. and Naranjo, L., An n -factor Gaussian model of oil futures prices. *J. Fut. Mkts*, 2006, **26**, 243–268.
- Culp, C.L. and Miller, M.H., Metallgesellschaft and the economics of synthetic storage. *J. Appl. Corp. Finance*, 1995, **7**, 62–76.
- Gibson, R. and Schwartz, E.S., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Heath, D., Jarrow, R. and Morton, A., Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica*, 1992, **60**, 77–105.
- Korn, O., Drift matters: an analysis of commodity derivatives. *J. Fut. Mkts*, 2005, **25**, 211–241.
- Mello, A.S. and Parsons, J.E., Maturity structure of a hedge matters: lessons from the Metallgesellschaft debacle. *J. Appl. Corp. Finance*, 1995, **8**, 106–121.
- Miltersen, K.R. and Schwartz, E.S., Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates. *J. Financial Quant. Anal.*, 1998, **33**, 33–59.
- Neuberger, A., Hedging long-term exposures with multiple short-term futures contracts?. *Rev. Financial Stud.*, 1999, **12**, 429–459.
- Schwartz, E.S., The stochastic behavior of commodity prices: implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Schwartz, E.S. and Smith, J.E., Short-term variations long-term dynamics in commodity prices. *Mgmt Sci.*, 2000, **46**, 893–911.

Appendix A: Expectation and Covariance Matrix of State Variables

$$E_t \begin{bmatrix} x_T^1 \\ x_T^2 \\ x_T^3 \end{bmatrix} = \begin{bmatrix} \mu_{11}(x_t^1, x_t^2, x_t^3, T-t) \\ \mu_{22}(x_t^1, x_t^2, x_t^3, T-t) \\ \mu_{33}(x_t^1, x_t^2, x_t^3, T-t) \end{bmatrix} = \begin{bmatrix} e^{-\kappa(T-t)}x_t^1 + \frac{\kappa(e^{-\gamma(T-t)} - e^{-\kappa(T-t)})}{\kappa - \gamma}x_t^2 + \frac{\kappa(e^{-\beta(T-t)} - e^{-\kappa(T-t)})}{\kappa - \beta}x_t^3 + \frac{\alpha}{\beta} \left(1 - \frac{\kappa e^{-\beta(T-t)} - \beta e^{-\kappa(T-t)}}{\kappa - \beta} \right) \\ e^{-\gamma(T-t)}x_t^2 \\ e^{-\beta(T-t)}x_t^3 + \frac{\alpha}{\beta}(1 - e^{-\beta(T-t)}) \end{bmatrix}.$$

$$\text{Cov}_t \begin{bmatrix} x_T^1 \\ x_T^2 \\ x_T^3 \end{bmatrix} = \begin{bmatrix} \Sigma_{11}(T-t) \Sigma_{12}(T-t) \Sigma_{13}(T-t) \\ \Sigma_{12}(T-t) \Sigma_{22}(T-t) \Sigma_{23}(T-t) \\ \Sigma_{13}(T-t) \Sigma_{23}(T-t) \Sigma_{33}(T-t) \end{bmatrix},$$

$$\begin{aligned} \Sigma_{11} = & \sigma_1^2 \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} + \sigma_2^2 \left(\frac{\kappa}{\kappa - \gamma} \right)^2 \left(\frac{1 - e^{-2\gamma(T-t)}}{2\gamma} + \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} - 2 \frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} \right) \\ & + \sigma_3^2 \left(\frac{\kappa}{\kappa - \beta} \right)^2 \left(\frac{1 - e^{-2\beta(T-t)}}{2\beta} + \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} - 2 \frac{1 - e^{-(\kappa+\beta)(T-t)}}{\kappa + \beta} \right) \\ & + 2\rho_{12}\sigma_1\sigma_2 \frac{\kappa}{\kappa - \gamma} \left(\frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} - \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} \right) + 2\rho_{23}\sigma_2\sigma_3 \frac{\kappa^2}{(\kappa - \gamma)(\kappa - \beta)} \left(\frac{1 - e^{-(\beta+\gamma)(T-t)}}{\beta + \gamma} \right. \\ & \left. - \frac{1 - e^{-(\kappa+\beta)(T-t)}}{\kappa + \beta} - \frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} + \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} \right) + 2\rho_{13}\sigma_1\sigma_3 \frac{\kappa}{\kappa - \beta} \left(\frac{1 - e^{-(\kappa+\beta)(T-t)}}{\kappa + \beta} - \frac{1 - e^{-2\kappa(T-t)}}{2\kappa} \right), \end{aligned}$$

$$\Sigma_{22} = \sigma_2^2 \frac{1 - e^{-2\gamma(T-t)}}{2\gamma}, \quad \Sigma_{33} = \sigma_3^2 \frac{1 - e^{-2\beta(T-t)}}{2\beta}, \quad \Sigma_{23} = \rho_{23}\sigma_2\sigma_3 \frac{1 - e^{-(\beta+\gamma)(T-t)}}{\beta + \gamma},$$

$$\begin{aligned} \Sigma_{12} = & \rho_{12}\sigma_1\sigma_2 \frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} + \sigma_2^2 \frac{\kappa}{\kappa - \gamma} \left(\frac{1 - e^{-2\gamma(T-t)}}{2\gamma} - \frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} \right) \\ & + \rho_{23}\sigma_2\sigma_3 \frac{\kappa}{\kappa - \beta} \left(\frac{1 - e^{-(\beta+\gamma)(T-t)}}{\beta + \gamma} - \frac{1 - e^{-(\kappa+\gamma)(T-t)}}{\kappa + \gamma} \right), \end{aligned}$$

$$\begin{aligned} \Sigma_{13} = & \rho_{13}\sigma_1\sigma_3 \frac{1 - e^{-(\kappa+\beta)(T-t)}}{\kappa + \beta} + \sigma_3^2 \frac{\kappa}{\kappa - \beta} \left(\frac{1 - e^{-2\beta(T-t)}}{2\beta} - \frac{1 - e^{-(\kappa+\beta)(T-t)}}{\beta + \kappa} \right) \\ & + \rho_{23}\sigma_2\sigma_3 \frac{\kappa}{\kappa - \gamma} \left(\frac{1 - e^{-(\beta+\gamma)(T-t)}}{\beta + \gamma} - \frac{1 - e^{-(\kappa+\beta)(T-t)}}{\kappa + \beta} \right). \end{aligned}$$

Appendix B: The Number of Futures Units

	WTI						Copper				
	1 Dec	2 Dec	3 Dec	4 Dec	5 Dec	6 Dec	1 Dec	2 Dec	3 Dec	4 Dec	5 Dec
1-2-3	1.20	-5.65	5.48	0	0	0	2.29	-8.12	6.84	0	0
1-2-4	0.62	-2.28	0	2.69	0	0	1.06	-2.89	0	2.83	0
1-2-5	0.35	-1.19	0	0	1.87	0	0.55	-1.34	0	0	1.80
1-2-6	0.20	-0.67	0	0	0	1.49	-	-	-	-	-
1-3-4	0.22	0	-3.70	4.50	0	0	0.39	0	-3.76	4.38	0
1-3-5	0.12	0.00	-1.47	0.00	2.37	0	0.20	0	-1.35	0	2.15
1-3-6	0.07	0	-0.74	0	0	1.69	-	-	-	-	-
1-4-5	0.06	0	0	-2.96	3.93	0	0.10	0	0	-2.44	3.35
1-4-6	0.03	0	0	-1.12	0	2.11	-	-	-	-	-
1-5-6	0.01	0	0	0	-2.38	3.39	-	-	-	-	-
2-3-4	0	1.29	-5.79	5.53	0	0	0	1.64	-5.90	5.25	0
2-3-5	0	0.64	-2.25	0	2.64	0	0	0.78	-2.13	0	2.36
2-3-6	0	0.34	-1.11	0	0	1.79	-	-	-	-	-
2-4-5	0	0.23	0	-3.52	4.32	0	0	0.29	0	-2.97	3.69
2-4-6	0	0.12	0	-1.31	0	2.21	-	-	-	-	-
2-5-6	0	0.05	0	0	-2.57	3.53	-	-	-	-	-
3-4-5	0	0	1.23	-5.45	5.24	0	0	0	1.27	-4.74	4.48
3-4-6	0	0	0.58	-1.99	0	2.43	-	-	-	-	-
3-5-6	0	0	0.20	0	-3.02	3.84	-	-	-	-	-
4-5-6	0	0	0	1.04	-4.59	4.57	-	-	-	-	-

An Empirical Study of the Impact of Skewness and Kurtosis on Hedging Decisions

4.1	Introduction	53
4.2	The Hedging Model	56
4.3	The Econometrics Model	58
	The S_U -normal Distribution • An Asset Return Model	
4.4	Empirical Results	60
	Data and Preliminary Analysis • Econometric Model Selection • Modelling Hedge Ratios • Hedged Portfolio Characteristics	
4.5	Conclusion	67
	References	67

Jing-Yi Lai

This study uses real price data rather than a simulation approach to investigate how hedging behaviours may change when hedgers consider skewness and excess kurtosis of hedging returns in their decision models. The study involves modelling the time-varying skewness and excess kurtosis of returns. The empirical results show that adding a preference for positively skewed returns to traditional mean–variance models may not lead to more speculative hedging/investment behaviours. Post-hedged return distributions suggest that the third moments of hedged portfolios have probably been well adjusted by mean–variance strategies, rendering three-moment decision models on a par with traditional mean–variance models. Additionally, considering the aversion to excess kurtosis will cause investors to hedge more. The research also provides empirical support for traditional minimum-variance strategies.

Keywords: Futures hedging; Skewness and kurtosis; S_U -normal distribution

4.1 Introduction

Following the landmark article of Johnson (1960), a large body of literature concerning futures hedging has centred on the minimum-variance (MV) strategy. The objective of variance minimisation assumes a high degree of risk aversion, so the MV hedging strategy operates as agents are motivated to minimise asset value risks by choosing an optimal hedge ratio, i.e. the size of futures contracts used to cover a cash position. Under the expected-utility-maximisation paradigm and with the utility function solely approximated by the mean and the variance of asset returns, the MV optimal hedging ratio can be derived with the ensuing formula, which is easy to understand and concise to apply. This attractive property of the MV strategy has been responsible for its popularity in hedging studies.

Nevertheless, the simplified two-moment utility function assumed for the decision process is plausible only if investors have a quadratic utility function and/or if the hedging returns are normally distributed. The quadratic utility function exhibits properties that do not make sense intuitively in practice, such as increasing absolute risk aversion and increasing relative risk aversion.* On the other hand, the normality assumption on the behaviour of the returns distribution has been found to be a restrictive assumption in practice. A large amount of empirical evidence suggests that many financial asset returns display significant leptokurtosis and skewness and, therefore, that they are not normal (e.g. Engle 1982, Bollerslev 1986, Hansen 1994, Theodossiou 1998, Harvey and Siddique 1999, Wang and Fawson 2001, Wang *et al.* 2001, etc.). Recent evidence from Chen *et al.* (2008) proposes a formal test on the joint normality of futures and their underlying spot returns. They document that the null of normality is rejected for all 25 contracts considered in their paper.

In response to the frequent findings of asymmetry and leptokurtosis in distributions of asset returns, econometric models applied to futures hedging studies have been improved to capture these features of tails. These features are captured by using suitable error distributions, such as the student's *t*, the skewed-*t*, the generalised error distribution (GED), and the stable distribution, among others (see, for example, Kavussanos and Nomikos 2000, Jin 2007, Park and Jei 2010).† Instead of defining different specific error distributions, the GARCH-jump model (Chan and Young 2006), which allows for a mixture of smooth volatility movement with abrupt changes or jumps, is also capable of accounting for the leptokurtosis in unconditional distribution. Although reshaping the error distribution as just documented may improve hedging effectiveness in terms of risk reduction, the benefit results exclusively from better estimates of second-moment price dynamics. In other words, how the skewness and excess kurtosis *per se* may influence agents' hedging decisions beyond the MV setting is still unanswered.

Aside from non-normally distributed returns, both theoretical and empirical studies have expressed the viewpoint that investors have not only an aversion to variance but also a preference for positive skewness (Kraus and Litzenberger 1976, Simkowitz and Beedles 1978, Conine and Tamarkin 1981, Kane 1982). Positive skewness means that a distribution has a long right tail, and hence, it is more likely for extremely high returns to occur than for their lower counterparts. That is, investors are facing an investment opportunity in which they have greater chances to capture upside potentials than to be trapped in a deep loss in return. This preference for a right-skewed portfolio is generally considered an inherent result of risk aversion (Scott and Horvath 1980). Excess kurtosis, which depicts fatter tails than in a normal distribution of asset returns, is almost universally observed in financial markets. It implies a greater likelihood that large values, both positive and negative, exist. In other words, facing excess-kurtosis returns, investors will have greater chances either to gain or lose a great deal. The very feature of tails sometimes is considered by investors in a similar way as variance, in that both aim to capture upside and downside deviations. Moreover, this kurtosis aversion, negative preferences for excess kurtosis, coincides with the restriction of decreasing absolute prudence (Kimball 1990) under the expected utility framework. Decreasing absolute prudence requires a negative fourth derivative of the expected utility function, which, coupled with decreasing absolute risk aversion, creates sufficient conditions for standard risk aversion (Kimball 1993, Dittmar 2002, Haas 2007).‡ For the above concerns, including an additional higher moment leads to a better approximation of the expected utility,

* Absolute risk aversion measures risk aversion for a given level of wealth, and relative risk aversion measures risk aversion for a given proportion of wealth. Casual empiricism suggests that risk aversion will probably decrease as one's wealth increases (Copeland *et al.* 2005).

† Brooks *et al.* (2002), Cotter and Hanly (2007) and Switzer and El-Khoury (2007) consider the effect of asymmetry on optimal hedging ratios. Rather than conditional skewness, the asymmetry here refers to differential impacts of positive and negative shocks on the conditional variance of the spot and futures markets. Their conditional distributions are specified to follow a normal skewness.

‡ Without this sufficient condition, some counterintuitive risk-taking behaviour may exist. For example, the agent may repeatedly accept a bet with negative expected payoff. Details of related discussions can be found in Pratt and Zeckhauser (1987) and Kimball (1993).

and for representative investors, particularly those who are not extremely risk averse, it seems reasonable to consider the third and fourth moments of returns in their utility function for hedging decisions.

The current literature points to a potential inadequacy in using traditional two-moment mean–variance models in hedging. In an early attempt to portray hedge decision making better, Karp (1987) formulates a hedging objective that explicitly incorporates higher moments of returns, whereas a more recent effort is found in Gilbert *et al.* (2006), Brooks *et al.* (2007) and Lien (2010). These studies, in general, suggest that skewness can increase the magnitude of speculative trading in forward markets under certain conditions. Gilbert *et al.* (2006) additionally show that skewness becomes a relevant concern as risk aversion decreases, contrary to the theoretical prediction from Karp (1987). With no parametric assumption on price movements, Brooks *et al.* (2007) suggest that incorporating higher moments into hedging decisions generates lower but higher-variable optimal hedging ratios than MV results, which indicates more active futures trading to be taken. Assuming that the spot price follows a skew-normal distribution, Lien (2010) finds that skewness, either positive or negative, promotes speculative trading in forward markets. The existing exposition approaches are implemented through theoretical discussion or numerical analysis under various hypothetical scenarios. In contrast, empirical investigations using real-world price data are sparse. The study in Zhang *et al.* (2009) is one of the very few exceptions. They report that the mean, skewness and utility of the hedging portfolio are improved from mean–variance to three-moment models. However, the empirical estimation of price behaviour is far from clear.

The simulation approach adopted in previous work has exceptional advantages in scenario analysis, which allows one to observe the exclusive effects induced by certain movements in price dynamics (e.g. skewness) while keeping other circumstances unchanged. However, this *ceteris paribus* assumption may not always hold in the real world. For example, skewness in equity returns may increase with its variances, as found in Lanne and Pentti (2007); a shift in the expected values of asset returns may accompany a shift in other moments (Brockett and Kahane 1992). The above examples suggest that the prediction based on hypothetical scenarios may not necessarily occur in practice.

Motivated by Gilbert *et al.* (2006), who consider skewness on top of the first two moments in their hedging decision model, this work further investigates how hedging behaviours may change when agents additionally consider excess kurtosis of hedging returns in their decision models. We compare the hedging/investment behaviours across models with different numbers of moments. The comparison involves sketching the degree of asymmetry and heaviness in the tails of returns from each hedging model. Using real price data instead of a simulation approach to depict the price movements, the current study may add to the literature by providing empirical evidence to compare higher moment hedging behaviour.

Our empirical work begins with specifying the first two moments of the Brent crude oil returns as a GARCH process while assuming the skewness and kurtosis parameters generated by an ARCH-type model. Our choice of the return error process is the S_U -normal distribution to be used in the maximum likelihood estimation (MLE) optimisation procedure for parameters of the hedging ratios. The S_U -normal distribution was first proposed by Johnson (1949), recently expanded by Choi and Nam (2008), and proven to be successful in describing the asymmetry and excess kurtosis of heteroscedastic financial returns. From the econometric estimation, this research recursively obtained the numerical values of return moments for each period over the hedging horizon. Using information available for the decision node in time, the hedging model with higher moments proposed in this study can be solved for optimal hedging ratios numerically, period by period.

The empirical results show that although the hedging/investment behaviours derived from mean–variance models are more speculative than three-moment models with additional skewness, the differences are not substantial in numerical scale. Further examination of the post-hedged return distributions suggests that third moments of the hedged portfolios have probably been well improved by mean–variance strategies, rendering three-moment decision models on a par with traditional mean–variance models. The empirical result is new to the literature. In addition, this paper provides evidence that adding the fourth moment to hedging models will lead to more conservative or less speculative hedging/investment behaviours compared to the mean–variance model and the three-moment model.

However, the disparity diminishes as the risk aversion increases. That is, for highly risk averse investors, differences in hedging strategies derived from the three objectives with different numbers of moments are insubstantial. The finding suggests that the traditional MV hedging ratios that completely ignore higher moments beyond the second of returns may still be credible.

This article is organised as follows. The next section describes the hedging model with higher moments and the derivation of optimal hedging ratios and then briefly discusses the S_U -normal distribution and the econometric model with time-varying moments. The empirical results will be provided, followed by concluding remarks in the final section.

4.2 The Hedging Model

Without loss of generality, this study assumes that an investor holds one unit of asset at current price p_0 at the initial period 0. To reduce the uncertainty in future asset value, the investor sells h unit(s) of forward or futures contracts written on the underlying asset at price f_0 . Here, h denotes the hedging ratio. By selling contracts, this agent will lock h unit(s) of asset at a value of f_0 ,^{*} leaving the remaining $(1 - h)$ unit(s) priced at the market value p_1 in the next period. In other words, if the investor chooses h to be 1, his or her future asset value is certain, essentially evaluated at f_0 . Highly risk averse investors, who may ignore a possible gain from a price increase in the spot market, will likely set $h = 1$. On the other hand, if h is chosen to be between zero and one, it is considered a normal hedging strategy in the sense that the agent sets part of his or her cash position to be protected by forwards or futures contracts and speculates an increase in price for the rest. If h is beyond the range of zero and one, then a pure speculating strategy is undertaken. When $h < 1$, the investor is short selling his cash stock. The value of the investor's hedged portfolio at time 1, $\tilde{\pi}_1$, can be expressed as[†]

$$\tilde{\pi}_1 = (1 - h)\tilde{p}_1 + f_0h = \tilde{p}_1 - h(\tilde{p}_1 - f_0). \quad (4.1)$$

The investment rate of return is capital gain divided by current asset value expressed in per cent as

$$\tilde{r}_\pi = \frac{[\tilde{p}_1 - h(\tilde{p}_1 - f_0)] - p_0}{p_0} \times 100. \quad (4.2)$$

Or equivalently,

$$\begin{aligned} \tilde{r}_\pi &= (1 - h)\tilde{r}_p + h \left[\frac{f_0}{p_0} - 1 \right] \times 100 \\ &= (1 - h)\tilde{r}_p + hr_b, \end{aligned} \quad (4.3)$$

where $\tilde{r}_p = [(\tilde{p}_1 - p_0)/p_0] \times 100$, which is the rate of return in the spot market, and

$$r_b = \left(\frac{f_0}{p_0} - 1 \right) \times 100,$$

which is known at time 0.

The current study further assumes that by choosing an optimal amount of futures contracts at period 0, the investor seeks to maximise the expected utility of the rate of returns on his or her hedged portfolio at period 1. Therefore, the expected-utility- maximisation argument can be written as

$$\max_h E[u(\tilde{r}_\pi)], \quad (4.4)$$

^{*} This study considers physical delivery in the commodity futures market.

[†] ' $\sim$ ' denotes that the value of the underlying variable is uncertain at the time the decision is made.

where $E[\cdot]$ is the expectation operator and $u(\cdot)$ is the utility function. Desirable properties for utility functions conventionally used in finance include $u'(\cdot) > 0$ for non-satiation, $u''(\cdot) < 0$ for risk aversion, and $u'''(\cdot) > 0$ for preference for positive skewness.* As this study intends to highlight the potential impact of the fourth moment on hedging decision making, we add a fourth condition, $u^{(4)}(\cdot) < 0$, to represent investors' dislike for excess kurtosis. A simple way to conceptualise this setting is to view excess kurtosis as the volatility of variance in returns. Alternatively, we may assume that the decision maker is resentful of a symmetrical increase in extreme returns.

To introduce moments higher than the second into the expected utility function, a straightforward technique is to take a Taylor series expansion of the utility function evaluated at the expected value of the elements, i.e. at $\tilde{r}_\pi = \bar{r}_\pi$.† Ignoring terms associated with moments higher than the fourth moment, the utility function can be rewritten by the Taylor series formula as

$$\begin{aligned} u(\tilde{r}_\pi) \approx & u(\bar{r}_\pi) + u'(\bar{r}_\pi) \cdot (\tilde{r}_\pi - \bar{r}_\pi) + \frac{1}{2!} u''(\bar{r}_\pi) \cdot (\tilde{r}_\pi - \bar{r}_\pi)^2 \\ & + \frac{1}{3!} u'''(\bar{r}_\pi) \cdot (\tilde{r}_\pi - \bar{r}_\pi)^3 + \frac{1}{4!} u^{(4)}(\bar{r}_\pi) \cdot (\tilde{r}_\pi - \bar{r}_\pi)^4. \end{aligned} \quad (4.5)$$

Taking the expectation of Equation (4.5), the objective function in (4.4) is transformed to

$$\begin{aligned} E[u(\tilde{r}_\pi)] = & u(\bar{r}_\pi) + \frac{1}{2} u''(\bar{r}_\pi) \cdot \sigma_\pi^2 + \frac{1}{6} u'''(\bar{r}_\pi) \cdot m_\pi^3 \\ & + \frac{1}{24} u^{(4)}(\bar{r}_\pi) \cdot m_\pi^4, \end{aligned} \quad (4.6)$$

where $\sigma_\pi^2 = E[(\tilde{r}_\pi - \bar{r}_\pi)^2]$, $m_\pi^3 = E[(\tilde{r}_\pi - \bar{r}_\pi)^3]$ and $m_\pi^4 = E[(\tilde{r}_\pi - \bar{r}_\pi)^4]$, representing the variance and the third and fourth central moments of the random variable $\tilde{r}_\pi$, respectively. To derive analytical solutions, it is necessary to specify an exact form of the utility function. Following Gilbert *et al.* (2006), a constant absolute risk aversion (CARA) utility function, also known as the exponential utility function, is considered‡

$$u(\tilde{r}_\pi) = -\exp(-\lambda \tilde{r}_\pi), \quad (4.7)$$

$$u'(\tilde{r}_\pi) = \lambda \exp(-\lambda \tilde{r}_\pi), u''(\tilde{r}_\pi) = -\lambda^2 \exp(-\lambda \tilde{r}_\pi), u'''(\tilde{r}_\pi) = \lambda^3 \exp(-\lambda \tilde{r}_\pi), \text{ and } u^{(4)}(\tilde{r}_\pi) = -\lambda^4 \exp(-\lambda \tilde{r}_\pi).$$

where λ is the coefficient of risk aversion and a higher λ represents a more risk-averse decision maker. The objective function hence becomes

$$E[u(\tilde{r}_\pi)] = -\exp(-\lambda \bar{r}_\pi) \left[1 + \frac{\lambda^2}{2} \sigma_\pi^2 + \frac{\lambda^3}{6} m_\pi^3 + \frac{\lambda^4}{24} m_\pi^4 \right]. \quad (4.8)$$

By the definition of r_π in Equation (4.2), the above objective function can be further simplified

$$\begin{aligned} E[u(\tilde{r}_p)] = & -\exp\{-\lambda[(1-h)\bar{r}_p + hr_b]\} \\ & \times \left[1 + \frac{\lambda^2}{2} (1-h)^2 \sigma_p^2 \right. \\ & \left. - \frac{\lambda^3}{6} (1-h)^3 m_p^3 + \frac{\lambda^4}{24} (1-h)^4 m_p^4 \right], \end{aligned} \quad (4.9)$$

* $u'(\cdot)$, $u''(\cdot)$ and $u'''(\cdot)$ are the first, second and third differentiations of the utility functions, respectively.

† Variables in the upper bar represent their expected values, i.e. $\bar{r}_\pi = E(\tilde{r}_\pi)$.

‡ The absolute risk-averse coefficient is defined as $-u''(\pi)/u'(\pi) = \lambda$, which is a constant. From Equation (4.7), one can easily obtain

where $\sigma_p^2 = E[(\tilde{r}_p - \bar{r}_p)^2]$, $m_p^3 = E[(\tilde{r}_p - \bar{r}_p)^3]$ and $m_p^4 = E[(\tilde{r}_p - \bar{r}_p)^4]$. A closed-form solution to the expected-utility-maximisation argument is given by the first derivative of Equation (4.9) with respect to the hedging ratio, h . The first-order condition is collected as

$$\begin{aligned} (1-h^*) - \frac{4}{\lambda} \left[\frac{1}{(\bar{r}_p - r_b)} + \frac{m_p^3}{m_p^4} \right] (1-h^*)^3 \\ + \frac{12}{\lambda^2} \left[\frac{m_p^3}{(\bar{r}_p - r_b)m_p^4} + \frac{\sigma_p^2}{m_p^4} \right] (1-h^*)^2 \\ - \frac{24}{\lambda^3} \left[\frac{\sigma_p^2}{(\bar{r}_p - r_b)m_p^4} \right] (1-h^*) + \frac{24}{\lambda^4 m_p^4} = 0, \end{aligned} \quad (4.10)$$

where h^* denotes the optimal value. In general, there may be as many as four unique values for h^* to solve Equation (4.10), while one value will maximise the expected utility function. This paper also obtains hedging performances under two alternative hedging objectives to compare with the four-moment hedging objective presented, where one accounts for the first three moments of a return distribution and the other ignores the skewness and kurtosis altogether. These two objectives maximise the following equations:

$$E[u(\tilde{r}_\pi)] = -\exp(-\lambda \bar{r}_\pi) \left[1 + \frac{\lambda^2}{2} \sigma_\pi^2 - \frac{\lambda^3}{6} m_\pi^3 \right]; \quad (4.11)$$

$$E[u(\tilde{r}_\pi)] = -\exp(-\lambda \bar{r}_\pi) \left[1 + \frac{\lambda^2}{2} \sigma_\pi^2 \right]. \quad (4.12)$$

The notations above have the same definition as in Equation (4.8). To facilitate our discussion, we refer to the three models of different hedging objectives (Equations 4.8, 4.11 and 4.12) as the four-moment, three-moment and mean-variance (alternatively two-moment) models, respectively.

4.3 The Econometrics Model

4.3.1 The S_U -normal Distribution

Among distributions that allow for time-varying skewness and excess kurtosis, the S_U -normal distribution is intuitive, and easy to understand and model.* If this distribution does not have restrictive assumptions for the range of parameters and has proved successful in describing the asymmetry and fat tails in heteroscedastic financial time series. The S_U -normal distribution is derived from a monotonic transformation of the normal distribution as follows.†

Consider a monotonic transformation of a random variable x from the standard normal variable z based on the following hyperbolic sine transformation of a normal variable‡:

$$x = \sinh(\delta + \theta z), \quad (4.13)$$

* Summarised by Choi and Nam (2008), distributions that allow for both the skewness and excess kurtosis property include the skewed- t by Hansen (1994), the non-central- t distribution by Harvey and Siddique (1999), the skewed GED by Theodossiou (1998) and the exponential generalised beta of the second kind (EGB2) by McDonald (1991, 1996), Wang and Fawson (2001) and Wang *et al.* (2001).

† More details of the derivation can be found in Choi and Nam (2008).

‡ By definition, $\sinh(y) = [\exp(y) - \exp(-y)]/2$.

where δ and θ are constant. Using the change-of-variable technique, the probability density function (p.d.f.) of the S_U -normal variable x is

$$f(x) = (2\pi)^{-1/2} J \exp\left(-\frac{1}{2}z^2\right), \quad (4.14)$$

where $J = \theta^{-1}(x^2 + 1)^{-1/2}$ is the Jacobian of the transformation and*

$$z = \theta^{-1} [\sinh^{-1}(x) - \delta]. \quad (4.15)$$

The shape parameters δ and θ in the p.d.f. determine the skewness and kurtosis of the distribution. The mean (m), variance (s^2), and the third and fourth central moments (μ_3 and μ_4) of x , obtained through the moment-generating function of the p.d.f., can be expressed as

$$m = \omega^{1/2} \sinh(\delta), \quad (4.16)$$

$$s^2 = \frac{1}{2}(\omega - 1) [\omega \cosh(2\delta) + 1] \mathbb{I}, \quad (4.17)$$

By definition, $\cosh(x) = [\exp(x) + \exp(-x)]$.

$$\mu_3 = \frac{1}{4} \omega^{1/2} (\omega - 1)^2 [\omega(\omega + 2) \sinh(3\delta) + 3 \sinh(\delta)], \quad (4.18)$$

$$\mu_4 = \frac{1}{8} (\omega - 1)^2 [\omega^2 (\omega^4 + 2\omega^3 + 3\omega^2 - 3) \cosh(4\delta) + 4\omega^2 (\omega + 2) \cosh(2\delta) + 3(2\omega + 1)], \quad (4.19)$$

where $\omega = \exp(\theta^2)$. It can be shown that the skewness coefficient of the S_U -normal distribution has the same sign as the value of δ . The greater the absolute value of δ , the more skewed the density will be. When $\delta = 0$, the S_U -normal distribution becomes symmetric. Moreover, the higher the value of θ , the more leptokurtic the density will be. As $\lambda = 0$ and θ is close to zero, the S_U -normal distribution has the same coefficients of skewness (= 0) and kurtosis (= 3).

4.3.2 An Asset Return Model

The only stochastic variable in the above hedging model is the rate of return in the spot market. The spot return (r_t) and its conditional variance (σ_t^2) are specified as a GARCH model such that

$$r_t = \mu + \varepsilon_t, \quad (4.20)$$

$$\sigma_t^2 = \beta_0 + \beta_1 \sigma_{t-1}^2 + \beta_2 \varepsilon_{t-1}^2, \quad (4.21)$$

where μ , β_0 , β_1 and β_2 are the coefficients being estimated and the error term $\varepsilon_t = \sigma_t u_t$ with u_t the innovation. The innovation term u_t is assumed to follow the S_U -normal distribution and can be written as $u_t = \alpha + \sinh(\delta + \theta z_t)$, where z_t is a standard normal variable.**† The location parameter α is included

* By definition, $\sinh^{-1}(x) = \ln(x + \sqrt{x^2 + 1})$.

† The mean, variance, and third and fourth central moments of ε_t , respectively, are

$$E(\varepsilon_t) = \sigma_t(\alpha + m), \quad (4.16')$$

$$\text{var}(\varepsilon_t) = \sigma_t^2 s^2, \quad (4.17')$$

$$m_3(\varepsilon_t) = \sigma_t^3 \mu_3, \quad (4.18')$$

$$m_4(\varepsilon_t) = \sigma_t^4 \mu_4, \quad (4.19')$$

to guarantee consistent quasi-maximum likelihood estimators (Newey and Steigerwald 1997, Choi and Nam 2008). Following the conditional variance dynamics specified in standard GARCH models, we assume that the parameters of skewness and kurtosis evolve over time in the following patterns:

$$\delta_t = c_0 + \sum_{i=1}^l c_i \delta_{t-i} + \sum_{j=1}^{\ell} w_j \left(\varepsilon_{t-j}^2 / \sigma_{t-j}^2 \right), \quad (4.22)$$

$$\theta_t = d_0 + \sum_{i=1}^q d_i \theta_{t-i} + \sum_{j=1}^p \bar{w}_j \left(\varepsilon_{t-j}^2 / \sigma_{t-j}^2 \right), \quad (4.23)$$

where m , s^2 , μ_3 and μ_4 are the same as in (4.16)–(4.19).

where $c_0, c_1, \dots, c_l, d_1, d_2, \dots, d_q; w_1, w_2, \dots, w_p; \bar{w}_1, \bar{w}_2, \dots, \bar{w}_p$ are coefficients. We then estimate the joint model in Equations (4.20) to (4.23) using the MLE method. The log likelihood function is given by

$$\begin{aligned} \ln L = & -\frac{T}{2} \ln(2\pi) \\ & -\frac{1}{2} \sum_{t=1}^T \left\{ \ln \theta_t^2 + \ln \sigma_t^2 + \ln \left[(\varepsilon_t \sigma_t^{-1} - a)^2 + 1 \right] \right\} \\ & -\frac{1}{2} \sum_{t=1}^T \theta_t^{-2} \left\{ \ln \left[(\varepsilon_t \sigma_t^{-1} - a) + \sqrt{(\varepsilon_t \sigma_t^{-1} - a)^2 + 1} \right] - \delta_t \right\}^2, \end{aligned} \quad (4.24)$$

where T is the number of observations.

4.4 Empirical Results

4.4.1 Data and Preliminary Analysis

The price data used in this study include daily closing prices for Brent crude oil (Brent blend) futures and its underlying spot prices traded on the Intercontinental Exchange (ICE). The ICE Brent crude oil futures was the second most actively traded contract in the worldwide commodity futures markets in 2006 and 2007 and remains second in the energy futures markets in 2008. The sample period is from 1 January 2000 to 31 December 2008. Both spot and futures prices are collected from CRB dataXtract, available from the Commodity Research Bureau. Table 4.1 summarises the descriptive statistics for the rate of returns (in per cent) of the Brent crude oil market. The Jarque–Bera normality test suggests that the unconditional distribution of spot returns is not normal and is possibly negatively skewed with fat tails. The augmented Dickey–Fuller unit root test confirms that the spot return series is stationary. Table 4.1 also presents the statistics for price levels in the oil spot and futures markets, as well as the fully hedged returns.*

4.4.2 Econometric Model Selection

Although the mean and variance Equations (4.20) and (4.21) of asset returns have been widely applied in empirical finance, the skewness and kurtosis parameter process in (4.22) and (4.23), respectively, is not as popular in the literature. To ensure specification of (4.22) and (4.23) best fitting the data, we try three alternatives before selecting the best one for hedge ratio estimation.

* Full-hedge strategy means that the cash position is completely protected by the futures contracts, leaving no uncertainty in future values. Mathematically, let $h = 1$ in Equation (4.2).

TABLE 4.1 Descriptive Statistics

	Spot Price	Futures Prices	Fully Hedged Returns	Spot Returns
Number of obs.	2306	2306	2306	2305
Mean	48.303	48.693	1.654	0.050
Standard dev.	27.340	27.026	4.683	2.296
Maximum	145.490	146.080	23.309	10.924
Minimum	16.570	17.680	-11.920	-20.040
Skewness coeff.	1.198	1.233	1.253***	-0.385***
Kurtosis coeff.	3.991	4.091	4.825***	7.751***
Jarque–Bera	645.560	698.379	923.035***	2224.531***
Unit root test				-50.608***

Notes: Fully hedged returns are defined as $(f_0 - p_0)/p_0 \times 100$, referring to Equation (4.2) when $h = 1$.

***Denotes that the null is rejected at the 1% significance level. The standard deviation of skewness and kurtosis coefficients are calculated as $\sqrt{6/T}$ and $\sqrt{24/T}$, where T is the sample size. Jarque–Bera is the statistic for normality testing. The unit root test is the augmented Dickey–Fuller test.

Model 1: δ and θ are constant.

Model 2:

$$\delta_t = c_0 + c\delta_{t-1} + w(\varepsilon_{t-1}^2 / \sigma_{t-1}^2), \quad (4.22a)$$

$$\theta_t = d_0 + d\theta_{t-1} + \varpi(\varepsilon_{t-1}^2 / \sigma_{t-1}^2). \quad (4.23a)$$

Model 3:

$$\delta_t = c_0 + \sum_{i=1}^{l'} c_i \delta_{t-i}, \quad (4.22b)$$

$$\theta_t = d_0 + \sum_{i=1}^{q'} d_i \theta_{t-i}. \quad (4.23b)$$

Model 1 restricts the shape parameters, δ and θ , to be constant, while Models 2 and 3 allow them to be time varying to different extents. The skewness and kurtosis equations in Model 2 are defined as a function of the first autoregressive term and the previous-period squared innovation $\varepsilon_{t-1}^2 / \sigma_{t-1}^2$. Model 3 assumes that the shape variables are exclusively related to their own lags. It may be acceptable here to term the specification as GARCH-type skewness and kurtosis for Model 2 and as ARCH-type skewness and kurtosis for Model 3.

Table 4.2 documents estimation results for the econometric models. The appropriate length of lags in the skewness and kurtosis equations of Model 3 are determined by the likelihood ratio tests. With few exceptions, almost all of the coefficients in the three models are highly significant. The Ljung–Box serial correlation tests show no evidence of serial correlation in either the standardised residuals ($\hat{z}$) or their squares. The Jarque–Bera normality test also fails to reject the null that the standardised residuals follow a normal distribution. Thus, the diagnostic statistics in Table 4.2 support the view that S_U -normal models fit the data well.

The skewness variable in Model 1 is significant at the 10% (5% against one-sided alternative) level, while the kurtosis variable is significant at the 1% level, suggesting that the model is successful in capturing the asymmetry and leptokurtosis of the data. The estimated coefficients of skewness and kurtosis in Model 1 can be obtained by applying Equations (4.16) to (4.19) and (4.16') to (4.19'), which yield -0.226 and 6.724, respectively. The likelihood ratio tests show that the goodness of fit significantly improves from Model 1 to either Model 2 or Model 3. This improvement implies that the degree of asymmetry and the heaviness in tails of the conditional distribution may not be constant.

TABLE 4.2 Results for the Econometric Model

Coefficient	Model 1	Model 2	Model 3
Conditional mean equation			
Constant	0.601 (0.159)***	2.201 (0.448)***	0.573 (0.157)***
α	-0.116 (0.079)	-0.649 (0.171)***	-0.082 (0.077)
Conditional variance equation			
Constant	0.090 (0.045)**	2.182 (0.760)***	0.065 (0.031)**
ε_{t-1}^2	0.066 (0.018)***	0.103 (0.033)***	0.059 (0.015)***
h_{t-1}	0.944 (0.013)***	0.606 (0.099)***	0.958 (0.009)***
Conditional skewness equation			
Constant	-0.074 (0.043)*	-0.095 (0.043)**	-0.105 (0.054)*
δ_{t-1}		0.423 (0.304)	-1.047 (0.062)***
δ_{t-2}			0.097 (0.053)*
δ_{t-3}			-0.294 (0.056)***
δ_{t-4}			-0.464 (0.047)***
δ_{t-5}			0.795 (0.057)***
δ_{t-6}			0.700 (0.071)***
$\varepsilon_{t-1}^2/h_{t-1}$		0.028 (0.009)***	
Conditional kurtosis equation			
Constant	0.655 (0.040)***	0.004 (0.002)*	0.038 (0.023)*
θ_{t-1}		0.982 (0.004)***	0.938 (0.037)***
$\varepsilon_{t-1}^2/h_{t-1}$		0.006 (0.001)***	
Log-likelihood	-4943.7	-4923.6	-4919.5
Likelihood ratio tests		40.061***	48.405***
$Q(12)$	9.203	10.248	8.886
$Q(24)$	27.324	25.385	25.373
$Q^2(12)$	10.481	9.691	9.485
$Q^2(24)$	22.486	23.818	26.236
Jarque-Bera	0.485	0.091	0.058
AIC	4.295	4.281	4.280
SC	4.298	4.287	4.288
HQ	4.292	4.277	4.274

Notes: Numbers in parentheses are the standard errors of the estimates.

***, **, * Denote that the test is significant at the 1%, 5% and 10% significance level, respectively. Likelihood ratio tests are conducted pairwise between Models 1 and 2, and Models 1 and 3. $Q(k)$ and $Q^2(k)$ are the Ljung-Box serial autocorrelation test on the standardised residual, $\hat{z}_t$, and its squares, where k is the number of the lag.

As the skewness and kurtosis variables are specified as time varying in Models 2 and 3, we report their average values over the sample period. The findings show that the average coefficients of skewness and kurtosis are -0.288 and 7.026, respectively, for Model 2, and -0.232 and 6.120, respectively, for Model 3. Model 2 has the smaller average skewness coefficient (more left-skewed) and the larger average kurtosis coefficient (heavier tails), but no significant differences exist between the two models. As Model 2 and Model 3 are not nested to one another, the likelihood ratio test is not an appropriate method to evaluate their goodness of fit. This research considers three alternative information criteria: the Akaike Information Criterion (AIC), the Schwarz Criterion (SC) and the Hanna-Quinn Criterion (HQC).^{*} The lower part of Table 4.2 shows that two of the information criteria (AIC and HQC) favour Model 3. The

* The information criteria are given by

hedging model in the following sections, therefore, assumes the form of Model 3, meaning that the investor frames price movement expectations in that particular manner.*

$$\begin{aligned} \text{AIC} &= -2(lT) + 2(k/T), \\ \text{SC} &= -2(l/T) + k \log(T)/T, \\ \text{HQC} &= -2(l/T) + 2k \log[\log(T)]/T, \end{aligned}$$

where l is the value of the log of the likelihood function with the k parameters estimated using T observations.

4.4.3 Modelling Hedge Ratios

Plugging the estimated moments into Equations (4.10) to (4.12), with λ set at a value of $\{100, 50, 10, 1, 0.75, 0.5, 0.1, 0.075, 0.05\}$ from the most to the least risk averse, the associated optimal hedging ratio can be solved period by period using price information up to the decision node. In detail, given the information at time $t - 1$ through Equations (4.16) to (4.19), we are able to calculate numerical values of the first to fourth moments of the innovation for the next period. The first four moments of error at t then can be obtained through Equations (4.16') to (4.19'). Finally, the forecasted values of r_t , σ_t^2 , δ_t and θ_t are yielded through Equations (4.20), (4.21), (4.22b) and (4.23b). With all these numerical values ready, the optimal hedging ratio h^* applied for period $t - 1$ to t can be reached by solving the four-degree polynomials of Equation (4.10).

Table 4.3 summarises the one-period-ahead optimal hedging ratios for various degrees of risk aversion from three different settings of hedging-objective models. At the top of each hedging model are the Means of Absolute values Deviating from unity (MAD, i.e. the average of $|h^* - 1|$). The unity hedging ratio in this study suggests a strategy in which the cash position is fully hedged, deviates from one, and is both positive and negative, implying that agents are involved in certain degrees of speculation. Furthermore, the larger the deviation, the more speculative the investment will be. Other summarised information includes the standard deviation of the optimal hedging ratios representing the variation in strategies, and the percentages of optimal hedging ratios across the sample period that are smaller than zero, larger than one and in between. As previously discussed herein, hedging ratios beyond zero and one represent pure speculating, whereas hedging ratios less than one imply a short sale on the spot market.

In general, optimal hedging ratios derived from the three hedging models of different return moments share similar patterns across various investor preferences for risks. Their MAD in the first row of each model in Table 4.3 clearly shows that, as decision makers' risk aversion increases, the MAD value approaches zero. Thus, more risk-averse investors choose an optimal hedging ratio closer to unity, suggesting a nearly full-hedge strategy. As expected, the standard deviation, representing the variation of the strategies, decreases with an increase in investors' degree of risk aversion. The percentage of optimal hedging ratios between zero and one provides similar information. For a risk aversion coefficient equal to or smaller than one, a short position in the cash market is considered, which is likely to happen when a sufficient amount of decrease in the spot price is speculated.

A careful comparison across the three hedging models indicates a variation in hedging behaviour. The results displayed at the top of Table 4.3 come from the objective, which considers all four moments of the hedged returns. Their hedging/investment behaviour, on average, seems to be the most conservative compared with the other two decision objectives, in that, given a degree of risk aversion, the four-moment model has the smallest MAD. The percentage of the four-moment model's hedging ratios within the range of zero and one is larger than those from the other two models in almost all cases. This

* Results for hedging behaviour from Models 1 and 2 are almost indistinguishable from those derived from Model 3, as will be presented later, and for which, to save space, we do not report. They are available upon request.

TABLE 4.3 Statistical Properties of the Optimal Hedging Ratios

Coefficient of risk aversion (λ)	100	50	10	1	0.75	0.5	0.1	0.075	0.05
(1) Four-moment hedging objective									
Mean of $ h^*-1 $	0.003**	0.005**	0.024**	0.243**	0.326**	0.489**	2.451**	3.245**	4.901**
Standard deviation of h^*	0.003**	0.006**	0.029**	0.288**	0.388**	0.584**	2.922**	3.853**	5.843**
$h^* < 0$ (%)	0.000**	0.000**	0.000**	0.044**	0.832**	5.035**	50.788**	53.940**	55.604**
$0 < h^* < 1$ (%)	59.019**	59.107**	58.888**	58.450**	58.450**	54.159**	8.363**	5.123**	3.546**
$h^* > 1$ (%)	40.981**	40.893**	41.112**	41.068**	40.718**	40.806**	40.849**	40.937**	40.849**
(2) Three-moment hedging objective									
Mean of $ h^*-1 $	0.003**	0.005**	0.027**	0.273**	0.363**	0.545**	2.743**	3.631**	5.485**
Standard deviation of h^*	0.004**	0.007**	0.034**	0.339**	0.451**	0.678**	3.463**	4.509**	6.925**
$h^* < 0$ (%)	0.000**	0.000**	0.000**	0.044**	1.883**	9.632**	36.559**	38.179**	39.580**
$0 < h^* < 1$ (%)	44.264**	44.264**	44.264**	44.221**	42.382**	34.632**	7.706**	6.086**	4.685**
$h^* > 1$ (%)	55.736**	55.736**	55.736**	55.736**	55.736**	55.736**	55.736**	55.736**	55.736**
(3) Two-moment (mean–variance) hedging objective									
Mean of $ h^*-1 $	0.003**	0.006**	0.028**	0.278**	0.370**	0.555**	2.775**	3.699**	5.549**
Standard deviation of h^*	0.004**	0.009**	0.034**	0.340**	0.453**	0.679**	3.396**	4.528**	6.792**
$h^* < 0$ (%)	0.000**	0.000**	0.000**	0.000**	1.051**	7.750**	36.602**	38.310**	39.580**
$0 < h^* < 1$ (%)	43.082**	43.082**	43.082**	43.082**	42.032**	35.333**	6.480**	4.772**	3.503**
$h^* > 1$ (%)	56.918**	56.918**	56.918**	56.918**	56.918**	56.918**	56.918**	56.918**	56.918**

Notes: h^* is the optimal hedging ratio. The statistical significance tests on the properties of the hedging ratios across models are conducted using pairwise comparison of the stationary bootstrap method of Politis and Romano (1994).

**, * Denote, respectively, that the differences in the given statistics are significant at the 5% and 10% levels under one-sided hypothetical tests. The pairwise comparisons are between the hedging ratios derived from the four-moment and the three-moment models; and from the three-moment and the mean–variance models.

result implies that the tendency of people to avoid extreme values in returns (kurtosis aversion) leads them to increase protection for their assets. The four-moment model has the smallest standard deviation of the optimal hedging ratios, especially those from highly risk-averse decisions. In other words, adding the aversion to excess kurtosis seems to create differences in hedging/investing activities, and the underscored differences become obvious for scant risk-averse investors.

To mitigate the concern that the numerical differences of hedging ratios resulting from different models are not statistically meaningful, a formal test on the significance is given. We follow Patton (2004) to adopt the stationary bootstrap resampling method of Politis and Romano (1994) for pairwise comparisons. The pairwise comparisons are conducted by investigating the bootstrap confidence interval on the differences of predefined statistics of interest, e.g., in our case, the MAD from two alternative hedging-objective models. Specifically, let the MAD of the hedging ratios from model i be μ_i , and that of model j be μ_j . Pseudo-time series are generated by resampling blocks of random length drawn from a

geometric distribution with the starting point of the blocks assumed to follow a discrete uniform distribution.* This procedure is repeated 1000 times to generate an approximate sampling distribution of the difference in MAD. To test the significance, if the lower bound of the bootstrap 5% confidence interval of $\mu_i - \mu_j$ is greater than zero, then we take μ_i to be significantly higher than μ_j at the 5% level in one-sided hypothetical testing. If, in another case, the upper bound of the interval is less than zero, then we take to be significantly higher than μ_j . If, on the other hand, the confidence interval includes zero, then the test suggests no evidence to reject the null that the two statistics are equal. The test first examines the differences in hedging ratios from various hedging objectives, and the results are summarised in Table 4.3. The pairwise comparison between the four-moment and the three-moment models shows that the MAD from the four-moment model is significantly smaller than that from the three-moment model at the 5% level. The variation of hedging ratios from the four-moment model is also significantly lower at the same level.

The MAD for the three-moment optimal hedging ratios is slightly smaller than for that in the two-moment (mean–variance) objective. Similarly, the percentages between zero and one are slightly higher in the three-moment model. Both results indicate that investors with an aversion to negative skewness in returns hedge more than those who are not concerned with skewness. Preferences for positive skewness do not lead to more speculative strategies regardless of how risk averse they are. Nevertheless, the differences between these two models are negligibly small, particularly for highly risk-averse agents. The pairwise comparison again is undertaken to test the significance of differences in hedging results between the models. The testing results show that the differences in the MAD of the hedging ratios between the two models are significant at the 10% level but fail to reject the null at the 5% level. No evidence is found to support the existence of differences in the variation of hedging ratios at even the 10% level. This finding suggests that considering skewness in portfolio returns to the traditional mean–variance models has no substantial impact on hedging/investment decisions. The underlying results using our oil price data are hence, quite different from previous work, e.g. Karp (1987) and Gilbert *et al.* (2006).

In examining the results across the three models, the disparity of hedging strategies is found to diminish as the risk aversion increases. That is, strategies taken by highly risk-averse investors are similar no matter which objective is considered. No substantial difference is observed in these risk categories. This empirical finding further supports the traditional MV hedging strategies that are based on the simple mean–variance framework. Ignoring higher moments of returns, especially skewness and kurtosis, may still conclude the same optimal hedging behaviour under an extremely risk-averse assumption.

4.4.4 Hedged Portfolio Characteristics

One possible reason for the negligibly dissimilar findings between the two- and three-moment models may be that for hedging/investment actions based on the two-moment (mean–variance) model, the degree of asymmetry in portfolio returns has been well modified. The statistics for the hedged portfolios in Table 4.4 support this explanation. Table 4.4 summarises the mean, standard deviation, and skewness and kurtosis coefficients of the hedged portfolio returns. Comparable statistics for the before-hedge asset returns and the fully hedged returns are available in Table 4.1. The lower part of Table 4.4 shows that the skewness coefficients for the hedged portfolios from the two-moment (mean–variance) model are all positive† and have similar values to those from the three-moment model given any degree of risk aversion except for nearly risk neutral decisions ($\lambda = 0.05$). This finding confirms that the third moments of the hedged portfolios may have been well adjusted by the strategies from the two-moment (mean–variance) model. Adding the preference for positive skewness to the model, therefore, has a limited impact.

* To apply the stationary bootstrap method of Politis and Romano (1994), smoothing parameter q is set to 0.9. The process of block resampling is stopped once the number of observations in the pseudo-time series reaches the sample size of our testing period, which is counted as one simulation.

† The coefficient of skewness of the unhedged portfolio returns found in Table 4.1 is -0.3849 .

TABLE 4.4 Statistics of the After-Hedge Portfolio Returns

Coefficient of risk aversion (λ)	100	50	10	3	0.75	0.5	0.1	0.075	0.05
(1) Four-moment hedging objective									
Mean	1.626**	1.623**	1.605**	1.549**	1.310**	1.154**	-0.677**	-1.534**	-2.982**
Standard deviation	5.095**	5.090**	5.047**	4.923**	4.517**	4.366**	9.173**	12.220**	19.254**
Coefficient of skewness	0.986**	0.983**	0.961**	0.891**	0.587**	0.435**	0.708**	0.243**	0.104**
Coefficient of kurtosis	4.687**	4.681**	4.641**	4.527**	4.400**	4.913**	9.821**	7.842**	7.622**
(2) Three-moment hedging objective									
Mean	1.631**	1.633**	1.652**	1.707**	1.944**	2.104**	4.002**	4.788**	6.375**
Standard deviation	5.102**	5.104**	5.115**	5.153**	5.398**	5.630**	10.611**	13.183**	18.572**
Coefficient of skewness	0.990**	0.991**	1.000**	1.024**	1.061**	1.034**	0.245**	0.099**	-0.042**
Coefficient of kurtosis	4.692**	4.691**	4.689**	4.684**	4.649**	4.615**	5.243**	5.629**	6.145**
(3) Two-moment (mean–variance) hedging objective									
Mean	1.631**	1.634**	1.657**	1.723**	2.007**	2.197**	4.471**	5.418**	7.313**
Standard deviation	5.103**	5.106**	5.126**	5.191**	5.543**	5.837**	11.228**	13.879**	19.379**
Coefficient of skewness	0.990**	0.992**	1.004**	1.036**	1.100**	1.089**	0.420**	0.280**	0.136**
Coefficient of kurtosis	4.692**	4.693**	4.697**	4.705**	4.698**	4.658**	4.593**	4.792**	5.133**

Notes: Given the optimal hedging ratio, h^* , for each period, the after-hedge portfolio return is calculated as in Equation (4.3). The statistical significance tests on the descriptive statistics of the portfolio returns are conducted using pairwise comparison of the stationary bootstrap method of Politis and Romano (1994).

**, * Denote, respectively, that the differences in the given statistics are significant at the 5% and 10% levels under one-sided hypothetical tests. The pairwise comparisons are between the portfolios from: (1) the four-moment and the three-moment models; (2) the three-moment and the mean–variance models; and (3) the mean–variance and the unhedged portfolios.

The significance tests on the differences between unhedged and after-hedged portfolios are undertaken using the aforementioned stationary bootstrap resampling method of Politis and Romano (1994). The results are presented in the lower part of Table 4.4, denoted by an asterisk after the numerical results. These results show that, except for the kurtosis coefficients in certain cases, the changes of other indications are significant at the 5% level. Specifically, hedging activities from the mean–variance strategies significantly increase portfolio returns, skew the return distribution to the right, and generate lower extreme values while in trading for a higher return volatility.

Table 4.4 also shows that the results from the two-moment and three-moment models across various risk preferences are seemingly more consistent with our intuition, in that highly risk-averse strategies lead to lower average returns with lower variations. The trade-off relationship between returns and risks still holds for the four-moment model, though in a counterintuitive direction when $\lambda > 0.5$. The more risk averse, the higher the average returns and the return variation will be. The investment behaviour for $\lambda < 0.5$ becomes unpredictable. This unpredictability suggests that a subtle effect of excess kurtosis on hedging decisions may exist.

Since the four-moment model suggests the most conservative or the least speculative investment strategy as presented earlier, hedged portfolios in the four-moment model have the lowest average returns with the least variation among the three hedging models. The differences are all significant at the 5%

level. Their skewness coefficients also show the lowest positive values. Compared with the three-moment model, the results show that adding the fourth moment to the objective function substantially reduces variations in hedged portfolio returns by trading for less positive-skewed portfolio returns. While the results between the three-moment and two-moment models do not seem to differ, they are still statistically significant. Thus, again, adding preference for positive skewness from the mean–variance setting will not have a substantial impact on the resulting hedged returns, with some of the differences on a decimal scale. However, they are not meaningless in a statistical sense.

4.5 Conclusion

Using real price data, this study investigates how hedging behaviour may change when hedgers consider higher moments of their hedged returns distribution, specifically for skewness and excess kurtosis. This work specifies a hedging decision model that considers up to the fourth moment of future returns, and it assumes that the returns follow an S_U -normal distribution to capture their conditional skewness and kurtosis. By modelling the oil returns in a GARCH process with ARCH-type skewness and kurtosis parameters, we demonstrate that the time-varying higher moments of returns can be successfully captured. All diagnostic checks indicate that the S_U -normal model fits the data fairly well.

By entering the expected values of the four moments, the current study numerically solves the hedging-decision models period by period. The empirical results show that hedging/investment behaviour derived from the four-moment model is more conservative or less speculative compared to the three-moment and mean–variance models. The four-moment model produces optimal hedge ratios with a mean close to unity and with the smallest standard deviation among the three models. This research differs from previous studies by finding that the investment behaviour for two-moment (mean–variance) models is more speculative than that for three-moment models, although the differences are negligibly small, particularly for highly risk-averse decisions. In other words, adding preference for positive skewness to traditional mean–variance models has no substantial impact on hedging/investment decisions. By cross investigating after-hedged returns, we claim that the above result is possibly caused by the consequence of the third moments of the hedged portfolios being properly modified by the mean–variance strategies. Therefore, adding the third moment does not have substantial impact on the optimal decision. Finally, as the differences in hedging strategies across objectives with different numbers of moments diminish for extremely risk-averse investors, we consider it to be empirical support for the traditional MV hedging models. Further empirical research may extend the study to the two stochastic prices model, allowing both futures and spot returns to change over time.

References

- Bollerslev, T., Generalized autoregressive conditional heteroskedasticity. *J. Econometr.*, 1986, **33**, 307–327.
- Brockett, P.L. and Kahane, Y., Risk, return, skewness and preference. *Mgmt. Sci.*, 1992, **38**, 851–866.
- Brooks, C., Cerny, A. and Miffre, J., Optimal hedging with higher moments. Mimeo, ICMA Centre, University of Reading, Reading, UK, 2007.
- Brooks, C., Henry, O.T. and Persaud, G., The effect of asymmetries on optimal hedge ratios. *J. Bus.*, 2002, **75**, 333–352.
- Chan, W.H. and Young, D., Jump hedges: An examination of movements in copper spot and futures markets. *J. Futures Mkts*, 2006, **26**, 169–188.
- Chen, S.S., Lee, C.F. and Shrestha, K., Do the pure martingale and joint normality hypotheses hold for futures contracts? Implications for the optimal hedge ratios. *Quart. Rev. Econ. Finan.*, 2008, **48**, 153–174.
- Choi, P. and Nam, K., Asymmetric and leptokurtic distribution for heteroscedastic asset returns: The S_U -normal distribution. *J. Emp. Finan.*, 2008, **15**, 41–63.

- Conine, T.E. and Tamarkin, M.J., On diversification given asymmetry in returns. *J. Finan.*, 1981, **36**, 1143–1155.
- Copeland, T.E., Weston, J.F. and Shastri, K., *Financial Theory and Corporate Policy*, 2005 (Pearson Press: Upper Saddle River, NJ).
- Cotter, J. and Hanly, J., Hedging effectiveness under conditions of asymmetry. MPRA Paper No. 3501, 2007. Available online at: <http://mpra.ub.uni-muenchen.de/3501/>
- Dittmar, R., Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *J. Finan.*, 2002, **57**, 369–403.
- Engle, R.F., Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 1982, **50**, 987–1007.
- Gilbert, S., Jones, S.K. and Morris, G.H., The impact of skewness in the hedging decision. *J. Futures Mkts*, 2006, **26**, 503–520.
- Haas, M., Do investors dislike kurtosis? *Econ. Bull.*, 2007, **7**, 1–9.
- Hansen, B.E., Autoregressive conditional density estimation. *Int. Econ. Rev.*, 1994, **35**, 705–730.
- Harvey, C.R. and Siddique, A., Autoregressive conditional skewness. *J. Finan. & Quant. Anal.*, 1999, **34**, 465–487.
- Jin, H.J., Heavy-tailed behaviour of commodity price distribution and optimal hedging demand. *J. Risk & Insur.*, 2007, **27**, 863–881.
- Johnson, L.L., The theory of hedging and speculation in commodity futures. *Rev. Econ. Studies*, 1960, **27**, 139–151.
- Johnson, N.L., Systems of frequency curves generated by method of translation. *Biometrika*, 1949, **36**, 149–176.
- Kane, A., Skewness preference and portfolio choice. *J. Finan. & Quant. Anal.*, 1982, **17**, 15–25.
- Karp, L.S., Methods for selecting the optimal dynamic hedge when production is stochastic. *Amer. J. Agri. Econ.*, 1987, **69**, 647–657.
- Kavussanos, M.G. and Nomikos, N.K., Constant vs. time-varying hedging ratios and hedging efficiency in the BIFFEX market. *Transp. Res. Part E: Logist. & Transp.*, 2000, **36**, 229–248.
- Kimball, M.S., Precautionary saving in the small and in the large. *Econometrica*, 1990, **58**, 53–73.
- Kimball, M.S., Standard risk aversion. *Econometrica*, 1993, **61**, 589–611.
- Kraus, A. and Litzenberger, R.H., Skewness preference and the valuation of risk assets. *J. Finan.*, 1976, **31**, 1085–1100.
- Lanne, M. and Pentti, S., Modeling conditional skewness in stock returns. *Eur. J. Finan.*, 2007, **13**, 691–704.
- Lien, D., The effects of skewness on optimal production and hedging decisions: An application of the skew-normal distribution. *J. Futures Mkts*, 2010, **30**, 278–289.
- McDonald, J.B., Parametric models for partially adaptive estimation with skewed and leptokurtic residuals. *Econ. Lett.*, 1991, **37**, 273–278.
- McDonald, J.B., An application and comparison of some flexible parametric and semi-parametric qualitative response models. *Econ. Lett.*, 1996, **53**, 145–152.
- Newey, W.K. and Steigerwald, D.G., Asymptotic bias for quasi-maximum likelihood estimators in conditional heteroskedasticity models. *Econometrica*, 1997, **65**, 587–599.
- Park, S.Y. and Jei, S.Y., Estimation and hedging effectiveness of time-varying hedge ratios: Flexible bivariate GARCH approaches. *J. Futures Mkts*, 2010, **30**, 71–99.
- Patton, A.J., On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *J. Finan. Econometr.*, 2004, **2**, 130–168.
- Politis, N.D. and Romano, J.P., The stationary bootstrap. *J. Amer. Statist. Assoc.*, 1994, **89**, 1303–1313.
- Pratt, J. and Zeckhauser, R., Proper risk aversion. *Econometrica*, 1987, **55**, 143–154.
- Scott, R.C. and Horvath, P.A., On the direction of preference for moments of higher order than the variance. *J. Finan.*, 1980, **35**, 915–919.
- Simkowitz, M.A. and Beedles, W.L., Diversification in a three-moment world. *J. Finan. & Quant. Anal.*, 1978, **13**, 927–941.

- Switzer, L.N. and El-Khoury, M., Extreme volatility, speculative efficiency, and the hedging effectiveness of the oil futures markets. *J. Futures Mkts*, 2007, 27, 61–84.
- Theodossiou, P., Financial data and the skewed generalized t distribution. *Mgmt. Sci.*, 1998, 44, 1650–1661.
- Wang, K.L. and Fawson, C., Modeling Asian stock returns with a more general parametric GARCH specification. *J. Finan. Stud.*, 2001, 9, 23–52.
- Wang, K.L., Fawson, C., Barrett, C.B. and McDonald, J., A flexible parametric GARCH model with an application to exchange rates. *J. Appl. Econometr.*, 2001, 16, 521–536.
- Zhang, L.-B., Wang, C.-F. and Fang, Z.-M., Optimal hedging ratio model with skewness. *Syst. Engng*, 2009, 29, 1–6.

5

Long-Term Spread Option Valuation and Hedging

5.1	Introduction	72
5.2	Cointegrated Prices and Mean Reversion of the Spread	73
	Cointegration Tests • Risk Neutral and Market Measure Spread Mean Reversion Tests	
5.3	Modelling the Spread Process.....	76
	Spread Process in the Risk Neutral Measure • Spread Process in the Market Measure • Futures Pricing • Calibration	
5.4	Spread Option Pricing and Hedging	81
5.5	Crack Spread: Heating Oil/ WTI Crude Oil	82
	Data • Unit Root and Cointegration Tests • Model Calibration • Futures Spread Option Valuation	
5.6	Location Spread: Brent/WTI Crude Oil.....	87
	Data • Unit Root and Cointegration Tests • Model Calibration • Futures Spread Option Valuation	
5.7	Conclusion	90
	Acknowledgements	90
	References.....	90
	Appendix A	92

M.A.H. Dempster

Elena Medova

Ke Tang

This paper investigates the valuation and hedging of spread options on two commodity prices which in the long run are in dynamic equilibrium (i.e. cointegrated). The spread exhibits properties different from its two underlying commodity prices and should therefore be modelled directly. This approach offers significant advantages relative to the traditional two price methods since the correlation between two asset returns is notoriously hard to model. In this paper, we propose a two factor model for the spot spread and develop pricing and hedging formulae for options on spot and futures spreads. Two examples of spreads in energy markets – the *crack spread* between heating oil and WTI crude oil and the *location spread* between Brent blend and WTI crude oil – are analysed to illustrate the results.

Keywords: Commodity spreads; Spread options; Cointegration; Mean-reversion; Option pricing; Energy markets

JEL classification: G12

5.1 Introduction

Commodity spreads are important for both investors and manufacturers. For example, the price spread between heating oil and crude oil (*crack* spread) represents the value of production (including profit) for a refinery firm. If an oil refinery in Singapore can deliver its oil both to the US and the UK, then it possesses a real option of diversion which directly relates to the spread of WTI and Brent crude oil prices. There are four commonly used spreads: spreads between prices of the same commodity at two different locations (*location* spreads) or times (*calendar* spreads), between the prices of inputs and outputs (*production* spreads) or between the prices of different grades of the same commodity (*quality* spreads).*

A *spread option* is an option written on the difference (*spread*) of two underlying asset prices S_1 and S_2 , respectively. We consider European options with *payoff* the greater or lesser of $S_2(T) - S_1(T) - K$ and 0 at maturity T for *strike price* K and focus on spreads in the commodity (especially energy) markets (for both spot and futures). In pricing spread options it is natural to model the spread by modelling each asset price separately. Margrabe (1978) was the first to treat spread options and gave an analytical solution for strike price zero (the *exchange option*). Closed form valuation of a spread option is not available if the two underlying prices follow geometric Brownian motions (see Eydeland and Geman, 1998). Hence various numerical techniques have been proposed to price spread options, such as for example the Dempster and Hong (2000) fast Fourier transformation approach. Carmona and Durrleman (2003) offer a good review of spread option pricing.

Many researchers have modelled the spread using two underlying commodity spot prices (the two price method) in the unique *risk neutral measure* as[†]

$$\begin{aligned} dS_1 &= (r - \delta_1)S_1 dt + \sigma_{1,1}S_1 dW_{1,1}, \\ d\delta_1 &= k_1(\theta_1 - \delta_1) dt + \sigma_{1,2}dW_{1,2}, \\ dS_2 &= (r - \delta_2)S_2 dt + \sigma_{2,1}S_2 dW_{2,1}, \\ d\delta_2 &= k_2(\theta_2 - \delta_2) dt + \sigma_{2,2}dW_{2,2}, \end{aligned} \quad (5.1)$$

where S_1 and S_2 are the spot prices of the commodities and δ_1 and δ_2 are their convenience yields, and $W_{1,1}$, $W_{1,2}$, $W_{2,1}$ and $W_{2,2}$ are four correlated Wiener processes. This is the classical Gibson and Schwartz (1990) model for each commodity price in a complete market.[‡] The return correlation $\rho_{12} := E[dW_{1,1} dW_{2,1}]/dt$ plays a substantial role in valuing a spread option; trading a spread option is equivalent to trading the correlation between the two asset returns. However, Kirk (1995), Mbanefo (1997) and Alexander (1999) have suggested that return correlation is very volatile in energy markets. Thus assuming a *constant* correlation in (5.1) is inappropriate.

But there is another longer term relationship between two asset prices, termed *cointegration*, which has been little studied by asset pricing researchers. If a cointegration relationship exists between two

* For more details on these concepts see Geman (2005a).

[†] Boldface is used throughout to denote random entities – here conditional on S_1 and S_2 having realized values S_1 and S_2 at time t which is suppressed for simplicity of notation.

[‡] We adopt this model for three reasons. (1) It fits futures contract prices much better than the one factor mean-reverting log price model as shown by Schwartz (1997). (2) In examining the historical commodity prices used here, we show that WTI and Brent crude oil and heating oil prices are not mean-reverting. This has been found by many others, e.g. Girma and Paulson (1999) and Geman (2005b). (We also show that the spread is mean-reverting.) Since the Gibson–Schwartz model has a GBM backbone, we believe it matches historical commodity prices better. (3) Schwartz (1997) shows that futures volatility in the one factor model will decay to close to zero after ten years, but using the Enron dataset he shows that the volatility for market futures with maturities longer than 2 years fluctuates around 12%. However the two factor Gibson–Schwartz model can match the volatility term structure quite well.

asset prices the spread should be modelled *directly* over the long term horizon. Soronow and Morgan (2002) proposed a one factor mean reverting process to model the location spread directly, but do not explain under what conditions this is valid nor derive any results.* See also Geman (2005a) where diffusion models for various types of spread option are discussed.

In this paper, we use two factors to model the spot spread process and fit the futures spread term structure. Our main contributions are threefold. First, we give the first statement of the economic rationale for mean reversion of the spread process and support it statistically using standard cointegration tests on data. Second, the paper contains the first test of mean-reversion of *latent* spot spreads in both the risk neutral and market measures. Third, we give the first latent multi-factor model of the spread term structure which is calibrated using standard state-space techniques, i.e. Kalman filtering.

The paper is organized as follows. Section 5.2 gives a brief review of price cointegration together with the principal statistical tests for cointegration and the mean reversion of spreads. Section 5.3 proposes the two factor model for the underlying spot spread process and shows how to calibrate it. Section 5.4 presents option pricing and hedging formulae for options on spot and futures spreads. Sections 5.5 and 5.6 provide two examples in energy markets which illustrate the theoretical work and Section 5.7 concludes.

5.2 Cointegrated Prices and Mean Reversion of the Spread

A *spread* process is determined by the dynamic relationship between two underlying asset prices and the *correlation* of the corresponding returns time series is commonly understood and widely used. *Cointegration* is a method for treating the long run dynamic equilibrium relationships between two asset *prices* generated by market forces and behavioural rules. Engle and Granger (1987) formalized the idea of integrated variables sharing an equilibrium relation which turns out to be either stationary or to have a lower degree of integration than the original series. They used the term cointegration to signify co-movements among trending variables which could be exploited to test for the existence of equilibrium relationships within the framework of fully dynamic markets.

In general, the return correlation is important for short term price relationships and the price cointegration for their long run counterparts. If two asset prices are cointegrated (5.1) is only useful for *short term* valuation even when the correlation between their returns is known exactly. Since we wish to model long term spread we shall investigate the cointegration (long term equilibrium) relationship between asset prices. First we briefly explain the economic reasons why such a long run equilibrium exists between prices of the same commodity at two different locations, prices of inputs and outputs and prices of different grades of the same commodity.†

The *law of one price* (or *purchasing power parity*) implies that cointegration exists for prices of the same commodity at different locations. Due to market frictions (trading costs, shipping costs, etc.) the same good may have different prices but the mispricing cannot go beyond a threshold without allowing market arbitrages (Samuelson, 1964). Input (raw material) and output (product) prices should also be cointegrated because they directly determine supply and demand for manufacturing firms. There also exists an equilibrium involving a threshold between the prices of a *commodity* of different grades since they are substitutes for each other. Thus the spread between two spot commodity prices reflects the profits of producing (production spread), shipping (location spread) or switching (quality spread). If such long-term equilibria hold for these three pairs of prices, cointegration relationships should be detected in the empirical data.

* We are grateful to an anonymous referee for this reference.

† Calendar spreads can be modelled using the models for individual commodities such as the models proposed by Schwartz (1997) and Schwartz and Smith (2000). In this paper two different commodities are considered.

In empirical analysis economists usually use equations (5.2) and (5.3) to describe the cointegration relationship:

$$\mathbf{S}_{1t} = c_t + d\mathbf{S}_{2t} + \boldsymbol{\varepsilon}_t, \quad (5.2)$$

$$\boldsymbol{\varepsilon}_t - \boldsymbol{\varepsilon}_{t-1} = \omega\boldsymbol{\varepsilon}_{t-1} + \mathbf{u}_t, \quad (5.3)$$

where $\mathbf{S}_1$ and $\mathbf{S}_2$ are the two asset prices and $\mathbf{u}$ is a Gaussian disturbance. Engle and Granger (1987) demonstrate that the error term $\boldsymbol{\varepsilon}_t$ in (5.2) must be *mean reverting* (5.3) if cointegration exists. Thus the *Engle–Granger* two step test for cointegration directly tests whether ω is a significantly negative number using an augmented Dicky and Fuller (1979) test. Note that (5.2) can be seen as the dynamic equilibrium of an economic system. When the trending prices $\mathbf{S}_1$ and $\mathbf{S}_2$ deviate from the long run equilibrium relationship they will revert back to it in the future.

For both location and quality spreads $\mathbf{S}_1$ and $\mathbf{S}_2$ should ideally follow the *same* trend, i.e. d should be equal to 1.* Since gasoline and heating oil are cointegrated substitutes, the d value could be 1 for both the heating oil/crude oil spread and the heating oil/gasoline spread (Girma and Paulson, 1999). For our spreads of interest – production and location – d is treated here as 1.

Letting $\mathbf{x}_t$ denote the spread between two cointegrated spot prices $\mathbf{S}_1$ and $\mathbf{S}_2$ it follows from (5.2) and (5.3) in this case that

$$\mathbf{x}_t - \mathbf{x}_{t-1} = c_t - c_{t-1} - \omega(c_{t-1} - \mathbf{x}_{t-1}) + \mathbf{u}_t, \quad (5.4)$$

i.e. the spread of the two underlying assets is *mean reverting*. No matter what the nature of the underlying $\mathbf{S}_1$ and $\mathbf{S}_2$ processes,[†] the spread between them can behave quite differently from their individual behaviour. This suggests modelling the spread *directly* over a long run horizon because the cointegration relationship has a substantial influence in the long run. Such an approach gives at least three advantages over alternatives since it: (1) avoids modelling the correlation between the two asset returns; (2) catches the long run equilibrium relationship between the two asset prices; (3) yields an analytical solution for spread options (cf. Geman, 2005a).

5.2.1 Cointegration Tests

The Engle–Granger two step test is the most commonly used test for the cointegration of two time series. We first need to test whether each series generates a *unit root* time series. If two asset price processes are unit root but the spread process is not, there exists a cointegration relationship between the prices and the spread will not deviate outside economically determined bounds. The augmented *Dickey–Fuller* (ADF) test may be used to check for unit roots in the two asset price time series. The ADF test statistic uses an ordinary least squares (OLS) auto regression

$$\mathbf{S}_t - \mathbf{S}_{t-1} = \pi_0 - \pi_1 \mathbf{S}_{t-1} + \sum_{i=1}^p \pi_{i+1} (\mathbf{S}_{t-i} - \mathbf{S}_{t-i-1}) + \boldsymbol{\eta}_t, \quad (5.5)$$

to test for unit roots, where $\mathbf{S}_t$ is the asset price at time t , π_i , $i = 0, \dots, p$, are constants and $\boldsymbol{\eta}_t$ is a Gaussian disturbance. If the coefficient estimate of π_1 is negative and exceeds the critical value in Fuller (1976)

* However for production spreads such as the spark spread (the spread between the electricity price and the gas price) d may not be exactly 1. Usually 3/4 of a gas contract is equivalent to 1 electricity contract so that investors trade a 1 electricity/3/4 gas spread which represents the profit of electricity plants (Carmona and Durrleman, 2003).

[†] Especially for commodities where many issues have to be considered, such as jumps, seasonality, etc. Hence no commonly acceptable model exists for all commodities.

then the null-hypothesis that the series has no unit root is rejected. We can use an extension of (5.3) corresponding to (5.2) to test the cointegration relationship:

$$S_{1t} = c_t + dS_{2t} + \varepsilon_t, \quad (5.6)$$

$$\varepsilon_t - \varepsilon_{t-1} = \chi_0 + \chi_1 \varepsilon_{t-1} + \sum_{i=1}^p \chi_{i+1} (\varepsilon_{t-i} - \varepsilon_{t-i-1}) + u_t.$$

When estimate of χ_1 is significantly negative the hypothesis that cointegration exists between the two underlying asset price processes is accepted (Hamilton, 1994).

Another way to test for cointegration of two time series is the Johansen method, where the cointegration relationship is specified in the framework of an error-correcting *vector auto-regression* (VAR) process (see Johansen, 1991 for details). For the two examples of this paper we test the cointegration relationship using both the Engle–Granger and Johansen tests, which agree in both cases.

5.2.2 Risk Neutral and Market Measure Spread Mean Reversion Tests

Testing the mean reversion of the spot spread is a key step in testing cointegration of the two underlying price processes. However usually only futures spreads are observed in the market and spot spreads are *latent* (i.e. unobservable). The futures spread with fixed maturity date is a martingale in the risk neutral measure and is thus not mean reverting. If we assume a constant risk premium the futures spread with fixed maturity date is also not mean reverting in the market measure. It is therefore difficult to test empirically the mean reversion of the spot spread using futures markets data without options data. We nevertheless propose novel methods to test the mean reversion of the spot spread in both the risk neutral and market measures which do not require derivative prices.

We can use *ex ante* market data analysis to test whether risk neutral investors *expect* the future spread to revert in the *risk neutral measure*. This approach uses relations between spread *levels* and the *spread term structure slope* defined as the price change between the maturities of two futures spreads. A negative relationship between the spot spread level (or short term futures spread level) and the futures spread term structure slope shows that risk neutral investors expect mean reversion in the spot spread. Indeed, since each futures price equals the trading date expectation of the delivery date spot price in the risk neutral measure the current term structure of the futures spread reveals where risk neutral investors expect the spot spread to be at delivery.* To discover this negative relationship in the risk neutral measure we estimate the time series model

$$\chi_L - \chi_S = \zeta + \gamma \chi_S + \varepsilon, \quad (5.7)$$

where χ_L and χ_S are, respectively, long-end and short-end spread levels in the futures spread term structure and ε is a noise term. If the estimate of γ is significantly negative there is evidence that the spot spread is *ex ante* mean reverting in the risk neutral measure.

To detect mean reversion of the spot spread in the *market measure* we construct a time series using *historical* futures data which preserves the mean reversion of the latent spot spreads. In [Section 5.3](#) we will see that futures spreads with a *constant* time to maturity preserve the mean reversion of the spot spread price.† Thus to test mean reversion empirically we use an ADF test based on

* Bessembinder *et al.* (1995) attempt to discover *ex ante* mean reversion in commodity spot *prices* by this technique.

† Since usually the spot prices are not directly observable, we use futures with short maturity to represent the spot prices (cf. Clewlow and Strickland, 1999). Therefore, in performing an Engle–Granger two step test for mean reversion we use futures spread time series with short and constant time to maturity because they can represent both the latent spot spread and preserve its mean reversion.

$$\begin{aligned}
& F(t + \Delta t, t + \Delta t + \tau) - F(t, t + \tau) \\
&= \chi_0 + \chi_1 \cdot F(t, t + \tau) + \sum_{i=0}^p \chi_{i+2} [F(t - i\Delta t, t - i\Delta t + \tau) \\
&\quad - F(t - (i+1)\Delta t, t - (i+1)\Delta t + \tau)] + \varepsilon_{t+\Delta t},
\end{aligned} \tag{5.8}$$

where $F(t, t + \tau)$ is the *futures spread* of maturity $t + \tau$ observed at t , Δt is the sampling time interval and ε_t is a random disturbance. Note that $F(t, t + \tau)$ and $F(t + \Delta t, t + \Delta t + \tau)$ are spreads relating to different futures contracts. If the estimate of χ_1 is significantly negative then the spot spread is deemed to be mean reverting.

5.3 Modelling the Spread Process

We now present the two factor spot spread models in both the risk neutral and market measures.

5.3.1 Spread Process in the Risk Neutral Measure

In the risk neutral measure the underlying *spot spread* process x and the *long run factor* y satisfy

$$\begin{aligned}
dx_t &= k(\theta + \phi(t) + y_t - x_t)dt + \sigma dW, \\
dy_t &= -k_2 y_t dt + \sigma_2 dW_2, \\
EdWdW_2 &= \rho dt,
\end{aligned} \tag{5.9}$$

where x and y are two *latent* mean reverting factors with, respectively, *long run means* θ and 0 and *mean reversion speeds* k and k_2 and ϕ is a *seasonality* function specifying the seasonality in the spread.

The x factor represents the mean reverting spot spread and the y factor is also a mean reverting process with zero long run mean representing deviation from the long term equilibrium of the spot spreads. We can interpret the dynamics of the spot spread x_t as reverting to a *stochastic* long run mean $\theta + y_t$, which itself reverts (eventually) to θ . The seasonality function ϕ induced by the seasonality of the individual commodity prices is specified (cf. Durbin and Koopman, 1997; Richter and Sorensen, 2002) as

$$\phi(t) = \sum_{i=1}^K [\alpha_i \cos(2\pi it) + \beta_i \sin(2\pi it)], \tag{5.10}$$

where α_i and β_i are constants.

Many tests for cointegration assume the long run relationship between two commodity prices is constant during the period of study. Modelling the spread by one mean reverting x factor is consistent with such a relationship. However in reality the long run relationship between the underlying commodities can change due to inflation, economic crises, changes in consumers' behaviour, etc. Gregory and Hansen (1996) have attempted to identify structural breaks of the cointegration relationship. We assume the long run relationship changes continuously and adopt the second y factor to reflect these changes. Our model nests the one factor model by setting the y factor to be identically zero. In [Sections 5.5](#) and [5.6](#) we will compare the ability of the one and two factor models to fit observed futures spreads.

5.3.2 Spread Process in the Market Measure

To calibrate our model to market data we need a version in which risk is priced. We can incorporate risk premium processes for x and y in the drifts of our risk neutral models to return to the market measure. Previous studies assume *constant* risk premia when modelling *Ornstein–Uhlenbeck* (OU) processes (see,

e.g. Hull and White, 1990; Schwartz, 1997; Geman and Nguyen, 2005) and similarly we assume that the two factor model in the market measure satisfies

$$\begin{aligned} d\mathbf{x}_t &= [k(\theta + \varphi(t) + y_t - x_t) + \lambda] dt + \sigma d\mathbf{W}, \\ dy_t &= (-k_2 y_t + \lambda_2) \sigma_2 d\mathbf{W}_2, \\ E d\mathbf{W} d\mathbf{W}_2 &= \rho dt, \end{aligned} \quad (5.11)$$

where λ and λ_2 are the *risk premia* of the $\mathbf{x}$ and $\mathbf{y}$ processes, respectively.

With starting time ν and starting position x_ν , $\mathbf{x}_s$ and $\mathbf{y}_s$ at time s can be expressed as

$$\begin{aligned} \mathbf{x}_s &= x_\nu e^{-k(s-\nu)} + \left(\theta + \frac{\lambda}{k} \right) [1 - e^{-k(s-\nu)}] + \frac{y_\nu k}{k - k_2} [e^{-k_2(s-\nu)} - e^{-k(s-\nu)}] \\ &\quad + \frac{\lambda_2}{k_2} (1 - e^{-k\Delta t}) - \frac{\lambda_2 k}{k_2(k - k_2)} (e^{-k_2\Delta t} - e^{-k\Delta t}) + G(\nu, s) \\ &\quad + \frac{K\sigma_2}{k - k_2} \int_\nu^s [e^{-k_2(s-u)} - e^{-k(s-u)}] d\mathbf{W}_2(u) + \int_\nu^s e^{-k(s-t)} \sigma d\mathbf{W}(t), \end{aligned} \quad (5.12)$$

$$\mathbf{y}_s = y_\nu e^{-k_2(s-\nu)} + \int_\nu^s e^{-k_2(s-u)} \sigma_2 d\mathbf{W}_2(u), \quad (5.13)$$

where $G(\nu, s)$ denotes the *seasonality effect* given by

$$G(\nu, s) = \int_\nu^s k e^{-k(s-r)} \varphi(r) dr. \quad (5.14)$$

Then the standard deviation of $\mathbf{x}_s$ becomes

$$b_s = \sqrt{A_1 + A_2 + 2\rho A_3}, \quad (5.15)$$

where

$$\begin{aligned} A_1 &:= \frac{\sigma^2}{2k} [1 - e^{-2k(s-\nu)}], \\ A_2 &:= \left(\frac{1}{2k_2} [1 - e^{-2k_2(s-\nu)}] + \frac{1}{2k} [1 - e^{-2k_2(s-\nu)}] \right. \\ &\quad \left. - \frac{2}{(k+k_2)} [1 - e^{-(k_2+k)(s-\nu)}] \right) \cdot \frac{k^2 \sigma_2^2}{(k-k_2)^2}, \\ A_3 &:= \frac{k\sigma_2\sigma}{k-k_2} \left(\frac{1}{k+k_2} [1 - e^{-(k+k_2)(s-\nu)}] - \frac{1}{2k} [1 - e^{-2k_2(s-\nu)}] \right). \end{aligned}$$

$$\text{As } s \rightarrow \infty, b_s \rightarrow \sqrt{\frac{\sigma^2}{2k} + \frac{\sigma_2^2}{2k_2(1+k_2/k)} + \frac{\rho\sigma\sigma_2}{2(k+k_2)}}, \text{ a constant.}$$

On the other hand it is easy to show in the two price extended Gibson and Schwartz (1990) model (5.1) that both S_1 and S_2 are non-stationary (i.e. *not* mean reverting). One easy way to show this is to write (5.1) in the vector format, as for $i = 1, 2$,

$$\begin{bmatrix} d \ln S_i \\ d\delta_i \end{bmatrix} = \begin{bmatrix} r - \frac{1}{2}\sigma_{i,1}^2 \\ k_i\theta_i \end{bmatrix} dt + \Psi \begin{bmatrix} \ln S_i \\ \delta_i \end{bmatrix} dt + \begin{bmatrix} \sigma_{i,1}dW_{i,1} \\ \sigma_{i,2}dW_{i,2} \end{bmatrix}, \text{ where } \Psi: \begin{bmatrix} 0 & -1 \\ 0 & -k_i \end{bmatrix}.$$

As shown in Arnold (1974), only if the real parts of all the eigenvalues of Ψ are strictly negative is $\ln S_i$ stationary. However inspection of Ψ shows that one eigenvalue is zero, so that $\ln S_i$ is not stationary.

Moreover, the variances of both S_1 and S_2 increase to infinity with time. Thus the variance of the spread will also blow up asymptotically. This is not consistent with the behaviour of spreads between cointegrated commodity prices in historical market data (Villar and Joutz, 2006).

5.3.3 Futures Pricing

Define $F(t, T, x_t)$ as the *futures spread* (the spread of two futures prices) of maturity T observed in the market at time t when the *spot spread* is x_t . In the risk neutral measure the spot spread process x must satisfy the *no arbitrage condition*

$$E[X_T | X_t] = F(t, T, X_t), \quad (5.16)$$

i.e. in the absence of arbitrage the conditional expectation in the risk neutral measure of the spot spread at T with respect to the realized spot spread x_t at t is the futures spread observed at time $t < T$. This must hold because it is costless to enter a futures spread (long one future and short the other).

Thus, for the two factor model

$$\begin{aligned} F(t, T; x_t) &= x_t e^{-k(T-t)} + \theta[1 - e^{-k(T-t)}] \\ &\quad + \frac{\gamma_t k}{k - k_2} [e^{-k_2(T-t)} - e^{-k(T-t)}] + G(t, T). \end{aligned} \quad (5.17)$$

From Ito's lemma it follows that the *risk neutral* futures spread $F(t, T)$ process with fixed maturity date T satisfies

$$dF(t, T) = e^{-k(T-t)} \sigma dW + k\phi \sigma_2 dW_2, \quad (5.18)$$

where $\phi := [e^{-k_2(T-t)} - e^{-k(T-t)}] / (k - k_2)$. In the market measure (5.18) becomes

$$dF(t, T) = (\lambda e^{-k(T-t)} + k\lambda_2 \phi) dt + e^{-k(T-t)} \sigma dW + k\phi \sigma_2 dW_2. \quad (5.19)$$

The futures spread with *fixed* maturity date T following (5.19) is *not* mean reverting.

However, defining $\tau := T - t$ as the constant time to maturity, the futures spread (5.17) can be rewritten as

$$F(t, t + \tau) = x_t e^{-k\tau} + e^{-k(t+\tau)} \int_t^{t+\tau} (\theta + \varphi(u)) e^{ku} k du + \phi y_t. \quad (5.20)$$

Differentiating (5.20) with respect to t the process of futures spreads with a constant time to maturity τ satisfies

$$\begin{aligned}
dF(t, t+\tau) &= e^{-k\tau} d\mathbf{x}_t + k\phi d\mathbf{y}_t - k^2 e^{-k(t+\tau)} [\\
&\quad \times \int_t^{t+\tau} (\theta + \varphi(u)) e^{ku} du] dt + k e^{-k(t+\tau)} [\theta + \varphi(t)] \\
&\quad \times [e^{k(t+\tau)} - e^{kt}] dt.
\end{aligned} \tag{5.21}$$

Substituting for the integral from (5.20) and using (5.11)–(5.13) we obtain

$$\begin{aligned}
dF(t, t+\tau) &= k[\theta + \varphi(t) + y_t e^{-k\tau} + \frac{\lambda e^{-k\tau}}{k} + \phi \lambda_2 \\
&\quad - F(t, t+\tau)] dt + e^{-k\tau} \sigma dW + k\phi \sigma_2 dW_2.
\end{aligned} \tag{5.22}$$

Note that if we set $y_t := 0$, $\lambda_2 := 0$, $\sigma_2 := 0$ in (5.22), we obtain the analogous process for the one factor model. From (5.22) a futures spread with constant time to maturity is mean reverting with the *same* mean-reversion speed as the spot process in both the two factor model and its one factor restriction (with $y \equiv 0$). Note that when $\tau \rightarrow 0$ (5.22) converges to (5.11).

5.3.4 Calibration

A difficulty with the calibration of the two factor model is that the factors (or state variables) are not directly observable, i.e. they are latent. An approach to maximum likelihood estimation of the model is to pose the model in *state space* form and use the *Kalman filter* to estimate the latent variables. Harvey (1989) and Hamilton (1994) give good descriptions of estimation, testing and model selection for state space models.

The state space form consists of transition and measurement equations. The *transition equation* describes the dynamics of the underlying unobservable data-generating process. The *measurement equation* relates a multivariate vector of observable variables, in our case future prices of different maturities, to the unobservable state vector of *state variables*, in our case the $\mathbf{x}$ and $\mathbf{y}$ factors. The measurement equation is given by (5.17) with the addition of uncorrelated disturbances to take account of pricing errors. These errors can be caused by bid-ask spreads, non-simultaneity of the observations, etc.

More precisely, suppose the data are sampled at equally spaced times t_n , $n = 1, \dots, N$, with $\Delta t := t_{n+1} - t_n$ for the interval between two observations. Let $\mathbf{X}_n := [x_{t_n}, y_{t_n}]^T$ represent the vector of state variables at time t_n . We obtain the transition equation from the discretization of (5.11) in the form

$$\mathbf{X}_{n+1} = A\mathbf{X}_n + b_n + \mathbf{w}, \tag{5.23}$$

where $\mathbf{w}$ is a serially independent multivariate normal innovation with mean 0 and covariance matrix Q , and A , b_n and Q are given by

$$\begin{aligned}
A &:= \begin{bmatrix} e^{-k\Delta t} & \frac{k}{k-k_2}(e^{-k_2\Delta t} - e^{-k\Delta t}) \\ 0 & e^{-k_2\Delta t} \end{bmatrix}, \\
b_n &:= \begin{bmatrix} \frac{\lambda_2}{k_2}(1 - e^{-k\Delta t}) - \frac{\lambda_2 k}{k_2(k-k_2)}(e^{-k_2\Delta t} - e^{-k\Delta t}) \\ + \left(\theta + \frac{\lambda}{k}\right)(1 - e^{-k\Delta t}) + G(t_n, t_n + \Delta t) \\ \frac{\lambda_2}{k_2}(1 - e^{-k_2\Delta t}) \end{bmatrix},
\end{aligned}$$

$$Q := \begin{bmatrix} V_x & V_{xy} \\ V_{xy} & V_y \end{bmatrix}$$

with

$$\begin{aligned} V_x &:= \frac{1-e^{-2k\Delta t}}{2k} \sigma^2 + \left\{ \frac{1}{2k_2} [1-e^{-2k_2\Delta t}] + \frac{1}{2k} [1-e^{-2k\Delta t}] \right. \\ &\quad \left. - \frac{2}{k+k_2} [1-e^{-(k+k_2)\Delta t}] \right\} \frac{k^2 \sigma_2^2}{(k-k_2)^2} \\ &\quad + \frac{2\rho k \sigma_2 \sigma}{k-k_2} \left[\frac{1}{k+k_2} (1-e^{-(k+k_2)\Delta t}) - \frac{1}{2k} (1-e^{-2k\Delta t}) \right], \\ V_y &:= \sigma_2^2 \left(\frac{1-e^{-2k_2\Delta t}}{2k_2} \right. \\ &\quad \left. + \frac{k \sigma_2}{k-k_2} \left[\frac{1-e^{-2k_2\Delta t}}{2k_2} - \frac{1-e^{-(k+k_2)\Delta t}}{k+k_2} \right] + \frac{\rho(1-e^{-(k+k_2)\Delta t})}{k+k_2} \right). \end{aligned}$$

In preliminary study we found that the estimated correlation $\hat{\rho}$ between the long and short end fluctuations was insignificant. This makes economic sense in that long-end movements are slow and driven by fundamentals, while short-term movements are random, fast and driven by market trading activities. In the model calibration we thus assume the correlation ρ to be zero.*

Let

$$Z_n := [\ln(F(T_n, t_n + \tau_1)) \dots \ln(F(t_n, t_n + \tau_M))]^T,$$

where $\tau_1, \dots, \tau_M$ are *times to maturity* for 1, $\dots$, M futures contracts. Thus the measurement equation takes the form,

$$Z_n = C_n X_n + d_n + \varepsilon_n, \quad (5.24)$$

where

$$\begin{aligned} C_n &:= \begin{bmatrix} e^{-k\tau_1} & \frac{k}{k-k_2} (e^{-k_2\tau_1} - e^{-k\tau_1}) \\ \vdots & \vdots \\ e^{-k\tau_M} & \frac{k}{k=k_2} (e^{-k_2\tau_M} - e^{-k\tau_M}) \end{bmatrix}, \\ d_n &:= \begin{bmatrix} \theta(1-e^{-k\tau_1}) + G(t_n, t_n + \tau_1) \\ \vdots \\ \theta(1-e^{-k\tau_M}) + G(t_n, t_n + \tau_m) \end{bmatrix} \end{aligned}$$

* Assuming that innovations in the long and short run are uncorrelated has been used to analyse the long-run and short-run components of stock prices (e.g. Fama and French, 1988).

and ϵ_n is a serially independent multivariate normal innovation with mean 0 and covariance matrix $\xi^2 I_M$. To reduce the parameters to be estimated we assume the variance of the futures spread pricing errors of all maturities to be the same. This assumption is based on the practical requirement that our model should price the futures of different maturities equally well. We also assume that the futures spread pricing errors are independent across different maturities and their common variance is denoted by ξ^2 .

5.4 Spread Option Pricing and Hedging

If the underlying asset price follows a Gaussian process the European *call* price* with maturity T on this asset can be calculated as

$$c = B \frac{b_s}{\sqrt{2\pi}} \exp\left[-\frac{(a_s - K)^2}{2b_s^2}\right] + B(a_s - K) \Phi\left(\frac{a_s - K}{b_s}\right), \quad (5.25)$$

where B is the price of a discount bond, a_s and b_s are, respectively, the mean and standard deviation of the underlying at maturity, and K is the strike price of the option and Φ denotes the cumulative distribution function of the normal distribution. Its *delta* is given by

$$\Delta_c = \frac{\partial C}{\partial a_s} = B \Phi\left(\frac{a_s - K}{b_s}\right). \quad (5.26)$$

Since a spread can be seen as simultaneously long one asset and short the other the delta hedge yields an *equal volume hedge*, i.e. long and short the same value of commodity futures contracts. The spread distribution is Gaussian in our model so that (5.25) and (5.26), respectively, can be used to price and hedge spread options. Procedures to hedge long term options with short term futures are well established, see e.g. Brennan and Crewe (1997), Neuberger (1998) and Hilliard (1999).

If we price and hedge options on the *spot spread*, then $a_s := F(t, T)$ which can be calculated from (5.17) and b_s is given by (5.15). If we price the options on the *futures spread*, a_s is the market observed futures spread and

$$b_s := \sqrt{A_1^F + A_2^F + 2\rho A_3^F}, \quad (5.27)$$

where

$$\begin{aligned} A_1^F &:= \frac{\sigma^2}{2k} [e^{-2k(T-R)} - e^{-2k(T-t)}], \\ A_2^F &:= \frac{k^2 \sigma_2^2}{(k-k_2)^2} \left\{ \frac{1}{2k_2} [e^{-2k_2(T-R)} - e^{-2k_2(T-t)}] + \frac{1}{2k} [e^{-2k(T-R)} - e^{-2k(T-t)}] \right. \\ &\quad \left. - \frac{2}{(k+k_2)} [e^{-(k_2+k)(T-R)} - e^{-(k_2+k)(T-t)}] \right\}, \\ A_3^F &:= \frac{k \sigma \sigma_2}{k-k_2} \left\{ \frac{1}{k+k_2} [e^{-(k+k_2)(T-R)} - e^{-(k+k_2)(T-t)}] \right. \\ &\quad \left. - \frac{1}{2k} [e^{-2k(T-R)} - e^{-2k(T-t)}] \right\}, \end{aligned}$$

when the option maturity is R , the futures maturity is $T \geq R$ and the current time is t .

* We show here only call price, the put price can be obtained by using put-call parity.

Since this paper focuses on ‘long term’ option valuation and hedging, we shall now discuss the option maturities appropriate for using our valuation models. These depend on the mean-reversion speeds k and k_2 . The average *decay half-lives* $\ln 2/k$ and $\ln 2/k_2$ of the mean reverting spread process can be used to represent its mean-reversion strength. If the spread option maturity is longer than the average half decay time, our methodology is appropriate to value the option. We expect the traditional two price spread option model (5.1) to overvalue these longer term options because of the variance blow-up phenomenon discussed previously. Mbanefo (1997) noted that long-term (longer than 90 days) crack spread options will be overvalued if mean-reversion of the spreads is not considered. Before turning to examples we remark that jumps and stochastic volatility are not important for determining *long term* theoretical or empirical option prices (Bates, 1996; Pan, 2002) so that the present parsimonious model is appropriate for this purpose.

5.5 Crack Spread: Heating Oil/WTI Crude Oil

The *crack spread* between the prices of heating oil and WTI crude oil represents the gross revenue from refining heating oil from crude oil. In [Section 5.2](#) we saw that a deviation from the (equilibrium) input and output price relationship can exist for short periods of time, but a prolonged large deviation could lead to the production of more end products until the output and input prices are brought nearer their long-term equilibrium.

5.5.1 Data

The data for modelling the crack spread consist of NYMEX daily futures prices of WTI crude oil (CL) and heating oil (HO) from January 1984 to January 2005. The times to maturity of these futures range from 1 month to more than 2 years. In order to test for unit roots, a data point is collected each month by taking the price of the futures contract with one month (fixed) time to maturity. For example, if the trading day is 20 February 2000 then the futures contract taken for the time series is the 20 March 2000 maturity future. We also create a long-end crack spread with time to maturity 1 year. The methodology is exactly the same as with the 1 month time series, but due to data unavailability we only use data from January 1989 to January 2005 to construct the long-end crack spread.

To calibrate the two factor model we calculate monthly the futures spreads for five futures contracts from January 1989 to January 2005. The time step Δt is thus chosen to be 1 month and the futures contracts chosen have 1, 3, 6, 9 and 12 month times to maturity.

5.5.2 Unit Root and Cointegration Tests

[Figure 5.1](#) shows the 1 month futures prices of crude oil and heating oil. First, we conduct the ADF test on the individual heating and crude oil prices.

The first step of the Engle–Granger two step test ([Table 5.1](#)) does not reject the hypothesis that both crude oil and heating oil are unit-root time series, which is consistent with the empirical findings in Girma and Paulson (1999) and Alexander (1999). This also agrees with the results in Schwartz (1997) that the Gibson–Schwartz two factor model fits the oil data much better than the one factor mean reverting log price model (the first model in Schwartz, 1997) because the individual commodity prices are unlikely to be mean reverting (see also Routledge *et al.*, 2000).*

Next we effect the second step in the Engle–Granger test, i.e. test the spread time series by estimating (5.8). We find a very strong mean reverting speed significant at the 1% level, which suggests that cointegration does exist in the data. In other words, the mean-reversion of the crack spread does not appear to be caused by the separate mean-reversion of the heating and crude oil prices but by the long

* This is why we chose the Gibson–Schwartz model (5.1) as our benchmark.

run equilibrium (cointegration) between them. We also employ the Johansen cointegration test, which again shows a significant cointegration relationship.

To test the mean-reversion in the risk neutral measure, we estimate the regression (5.7) using the 1 year crack spread as the long-end futures spread and the 1 month crack spread as the short-end futures spread. The results are given in Table 5.2 for the 1 year and 1 month crack spreads depicted in Figure 5.2.

The estimate of γ is significantly less than zero. Thus both the market and risk neutral tests give evidence that the spot crack spread is mean reverting.




FIGURE 5.1 The 1 month futures prices of crude oil and heating oil.

TABLE 5.1 Cointegration Test for Crude Oil and Heating Oil

	No. of observations	Lag (p)	π_1 or χ_1	t value
<i>Engle-Granger two step tests</i>				
Crude oil	257	6	0.020	1.03
Heating oil	257	6	0.019	1.13
Crack spread	257	6	-0.31	-3.97*
<i>Johansen test</i>				
Likelihood ratio trace statistic: 46.81 (significant at 99%) ^a				

^aNote: that the null hypothesis is no cointegration, which is rejected at the 99% confidence level.

*Significant at the 1% level.

TABLE 5.2 Regression (5.7) Parameter Estimates for the Crack Spread

	S	γ
Value	2.16	-0.55
t -Stat.	15.69*	-16.39*
No. of observations	154	
R^2	57.37%	

*Significant at the 1% level.




FIGURE 5.2 The 1 month and 1 year crack spreads.

5.5.3 Model Calibration

We assume that the seasonality of the crack spread follows an *annual* pattern and that ϕ can be specified as

$$\varphi(t) = \alpha \cos(2\pi t) + \beta \sin(2\pi t).$$

Our simple method of dealing with seasonality aims to keep the model parsimonious. Thus the $G(v, s)$ function in (5.12) becomes

$$\begin{aligned} G(v, s) = & \frac{\alpha k^2}{k^2 + 4\pi^2} [\cos(2\pi s) - \cos(2\pi v)e^{-k(s-v)} + \frac{2\pi}{k} (\sin(2\pi s) \\ & - \sin(2\pi v)e^{-k(s-v)})] + \frac{\beta k^2}{k^2 + 4\pi^2} [\sin(2\pi s) - \sin(2\pi v)e^{-k(s-v)} \\ & - \frac{2\pi}{k} (\cos(2\pi s) - \cos(2\pi v)e^{-k(s-v)})]. \end{aligned}$$

Our calibration results are shown in Table 5.3. The market prices of risk λ and λ_2 are insignificantly different from zero in the two factor model,*but estimates of all the other parameters (α , β , σ , σ_2 , k , k_2 and θ) are significant. The estimated seasonality function ϕ is shown in Figure 5.3, where we can see that in the winter time the spread is more valuable than in summer. Since crude oil does not exhibit seasonality, the seasonality of the crack spread is caused by the seasonality of heating oil.†The *mean error* (ME) and the *root mean squared error* (RMSE) are plotted in Figure 5.4. The large pricing errors from 1990 to 1991 are due to the Gulf War.

The asymptotic estimate of the standard deviation of the spread is \$2.42. Also, assuming zero correlation between the two factors, we can examine the ratio of the short-end and long-end variance – A1:A2 in (5.15) – as time goes to infinity. It is about 2.5:1 in this example. To see whether the γ factor provides a statistical improvement, we can compare the differences in log-likelihood scores with and without

* Probably due to alternating periods of spread backwardation and contango implying corresponding periods of positive and negative market prices of risk due to shifts in input and output supply/demand balances (Geman, 2005a).

† Heating oil is known to be more expensive in winter than in summer due to high demand in the winter.

TABLE 5.3 Parameter Estimates for the Two Factor Model

	k	θ	λ	σ	k_2	λ_2	σ_2	α	β	ξ
Value	3.39	3.58	-1.33	5.32	0.66	0.25	1.63	1.61	-1.47	0.38
Std. error	0.13	0.05	1.44	0.08	0.27	0.41	0.13	0.08	0.18	0.06
Log likelihood	-771									

Note: We assume the correlation ρ between the two factors is zero.




FIGURE 5.3 The estimated seasonality function ϕ for the crack spread.




FIGURE 5.4 The mean pricing error (ME) and root mean squared error (RMSE) for the crack spread.

the y factor. The relevant statistic for this comparison is the chi-squared *likelihood ratio test* statistic (Hamilton, 1994) with three degrees of freedom and the 99th percentile of this distribution is 11.34. We thus ran the calibration with only the x factor and obtained its likelihood. Given that the likelihood ratio statistic is 282, the improvements provided by the y factor are highly significant. The correlation

between the filtered long and short factors is 0.0056, which is consistent with the assumption that the correlation between these two factors is zero.

5.5.4 Futures Spread Option Valuation

The average half-life is about 7 months for the x factor so that using the methodology of this paper is appropriate to valuing an option longer than 7 months.

For comparison it is necessary to calculate option values from a model which ignores the cointegration effect. Thus, we must simulate both the crude and heating oil futures prices and then calculate the spread option value* (a two price method). We utilize the Gibson and Schwartz (1990) model for each commodity. Appendix A gives a detailed justification of this method.

On 3 January 2005 the HO06N (heating oil future with maturity July 2006) contract had a value 44.2 (\$/Barrel) and implied volatility 29.9%; on the same day the CL06N (crude oil future with maturity July 2006) traded at 39.78 (\$/Barrel) and implied volatility 28.3%. Using the parameters in Tables 5.3–5.5 give, respectively, the European option values and deltas on the *futures* crack spread with maturity July 2006 on this date. As a comparison, we also list the option values and deltas for the one factor model (with $y \equiv 0$). Since as previously noted the correlation between energy price returns is difficult to obtain and very volatile, when simulating the two prices we here assume $\hat{\rho}_{12}$ to be the constant 20 year correlation between the crude and heating oil 1 month futures prices (0.89 in our calibration).

From Table 5.4 we can see that the option value from the one factor model is typically smaller than that from the two factor model and the latter is much smaller than that from the two price model. Since this latter model does not consider mean-reversion (the cointegration) of the spread, its spread distribution at maturity is wider than that of a cointegrated model and thus yields a larger option value. Put simply, a non-cointegrated model ignores the long run equilibrium between crude and heating oil prices and thus *overprices* the option. In Table 5.5, both the one factor and two factor models yield an equal volume hedge but the two price model does not. As is well known, the less disperse the underlying terminal distribution, the more sensitive are the option deltas to the strike prices.[†] Thus the one-factor model yields the most sensitive deltas and the two price model has the least sensitive deltas amongst the three models.

TABLE 5.4 Comparison of Crack Spread Option Values

Strike (\$)	2	3	4	5	6
One factor spread model	2.41	1.61	0.98	0.52	0.24
Two factor spread model	2.48	1.72	1.10	0.64	0.33
Two price model	3.87	3.32	2.83	2.39	2.03

TABLE 5.5 Comparison of Crack Spread Option Deltas

Strike (\$)	2		3		4		5		6	
Underlying	HO	CL	HO	CL	HO	CL	HO	CL	HO	CL
One factor spread model	0.84	−0.84	0.72	−0.72	0.55	−0.55	0.37	−0.37	0.21	−0.21
Two factor spread model	0.81	−0.81	0.69	−0.69	0.54	−0.54	0.39	−0.39	0.24	−0.24
Two price model	0.64	−0.56	0.57	−0.53	0.55	−0.43	0.43	−0.42	0.41	−0.36

HO, heating oil; CL, crude oil.

* Since there is no analytical solution when the strike is not zero a convenient way to calculate the option value is by Monte Carlo simulation.

[†] The *sensitivity* is defined as the ratio of the change of the deltas to the change of the strike prices.

5.6 Location Spread: Brent/WTI Crude Oil

We define the *location spread* as the price of WTI crude oil (CL) minus the price of the Brent blend crude oil (ITCO). WTI is delivered in the USA and Brent in the UK.

5.6.1 Data

NYMEX daily futures prices of WTI crude oil were described in the previous example. The daily Brent futures prices are from January 1993 to January 2005. The time to maturity of the Brent futures contracts range from 1 month to about 3 years. As in the previous example, monthly data is used to test for the unit root in Brent oil prices. We also create a monthly long end location spread with time to maturity of 1 year ranging from January 1993 to January 2005. In order to calibrate our model, we calculate monthly the futures spread with five maturities from January 1993 to January 2005. The five contracts involved are the 1, 3, 6, 9 and 12 month futures.

5.6.2 Unit Root and Cointegration Tests

As from the previous example we know that the WTI crude oil price follows a unit root process, in this example we need only conduct the ADF test on Brent crude oil prices. Figure 5.5 shows the 1 month futures prices of WTI crude oil and Brent blend.

Similar to WTI crude oil, the Brent blend price is also a unit root process, but the corresponding location spread appears to be a mean reverting process. This again suggests the existence of a long run equilibrium in the data. The Johansen test also confirms the existence of the cointegration relationship (see Table 5.6). The estimate of γ in (5.7) is strongly negative, so that the market appears to expect the spot spread to be mean reverting in the risk neutral measure (Table 5.7). The 1 year and 1 month spread evolution is depicted in Figure 5.6.

Hence both the market and risk neutral analyses support the mean reversion of the spot spread.




FIGURE 5.5 The 1 month futures prices of WTI and Brent crude oil.

TABLE 5.6 Cointegration Test for Brent and Wti Crude Oil

	No. of observations	Lag (p)	π_1 or χ_1	t -Stat
<i>Engle–Granger two step tests</i>				
Brent blend	152	6	0.04	2.48
WTI	257	6	0.02	1.03
Location spread	152	6	−0.28	−3.45*
<i>Johansen test</i>				
Likelihood ratio trace statistic: 26.74 (significant at 99%) ^a				


^a Note: that the null hypothesis is no cointegration, which is rejected at the 99% confidence level.

*Significant at the 1% level.

TABLE 5.7 Regression (5.7) Parameter Estimates for the Location Spread

	ζ	γ
Value	0.92	−0.70
t -Stat.	11.57*	−15.54*
No. of observations	152	
R^2	63.65%	

*Significant at the 1% level.

**FIGURE 5.6** The 1 month and 1 year location spreads.


5.6.3 Model Calibration

We did not find evidence of seasonality in the location spread. Thus, we set $\varphi(t) \equiv 0$. Table 5.8 lists the calibration results for our model. Figure 5.7 shows the pricing errors of the model. We see that the asymptotic standard deviation of the spread is estimated to be \$2.98. The ratio of short-end to the long-end variance ($A_1:A_2$ in (5.9)) is 1:5, i.e. the long end (second factor) movement of the spread accounts for

TABLE 5.8 Parameter Estimates of the Two Factor Model

	k	θ	λ	σ	k_2	λ_2	σ_2	ξ
Value	5.04	0.00	0.50	3.93	0.14	0.21	1.45	0.12
Std. error	0.28	0.00	1.18	0.28	0.02	0.28	0.10	0.03
Log likelihood	-23.8							

Note: We assume the correlation ρ between the two factors is zero.

**FIGURE 5.7** The mean pricing error (ME) and root mean squared error (RMSE) for the location spread.

much more variance than that of the short end (first factor). From the likelihood ratio test the two factor model is significantly better than the one factor model in explaining the observed LSBW spread data.*

5.6.4 Futures Spread Option Valuation

The average half-life is about 2.5 years, thus the methods in this paper should be used to price an option of maturity longer than 2.5 years.

On 1 December 2003 the ITCO06Z (Brent blend crude oil future with maturity December 2006) contract had a value 24.62 (\$/Barrel) and implied volatility 19.2%; on the same day the CL06Z (WTI crude oil future with maturity December 2006) had a value 25.69 (\$/Barrel) and implied volatility 19.2%. Tables 5.9 and 5.10, respectively, show the European option values and deltas on the futures spread with maturity December 2006 using the one and two factor model and the two price model. Note that $\hat{\rho}_{12}$ is estimated to be 0.94 for the 1 month WTI and Brent oil data. The correlation between the two factors is 0.04 which is again consistent with our zero correlation assumption.

Table 5.9 shows that, as in the previous example, the option value of the one factor model is typically smaller than that of the two factor model and the latter is much smaller than the two price model.

* The likelihood ratio statistic is 628.

TABLE 5.9 Comparison of Location Spread Option Values

Strike (\$)	−1	0	1	2	3
One factor spread model	1.96	1.19	0.60	0.24	0.08
Two factor spread model	2.11	1.43	0.89	0.49	0.24
Two price model	2.27	1.61	1.10	0.71	0.45

TABLE 5.10 Comparison of Location Spread Option Deltas

Strike (\$)	−1		0		1		2		3	
Underlying	CL	ITCO	CL	ITCO	CL	ITCO	CL	ITCO	CL	ITCO
One factor spread model	0.81	−0.81	0.68	−0.68	0.47	−0.47	0.26	−0.26	0.10	−0.10
Two factor spread model	0.74	−0.74	0.62	−0.62	0.47	−0.47	0.32	−0.32	0.19	−0.19
Two price model	0.69	−0.63	0.55	−0.53	0.48	−0.43	0.42	−0.34	0.34	−0.30

CL, WTI crude oil; ITCO, Brent crude oil.

As before, by ignoring cointegration the two price model tends to over value the long term option. We obtain a pattern of deltas similar to the previous example (see [Table 5.5](#)) and the explanation for this is the same.

5.7 Conclusion

In this paper we have developed spread option pricing models for the situation when the two underlying prices of the spread are cointegrated. Since the cointegration relationship is important for the long run dynamical relationship between the two prices, contingent claim evaluation based on spreads should take account of this relationship for long maturities. We model the spread process *directly* using a two factor model, i.e. we model directly the dynamic deviation from the long run equilibrium which cannot be specified correctly by modelling the two underlying assets separately. We also propose two methods (risk neutral and market) for testing data for mean-reversion of the spread process. The corresponding spread option can be priced and hedged analytically. In order to illustrate the theory we study two examples – of crack and location spreads, respectively. Both spread processes are found to be mean reverting. From likelihood ratio tests the second y factor is found to be important in explaining the crack and location spread data. The option values from our model are quite different from those of standard models, but they are consistent with the practical observations of Mbanefo (1997).

Acknowledgements

We are grateful to Helyette Geman and an anonymous referee for comments which materially improved the paper. This chapter is reprinted with permission from the Journal of Banking and Finance, 2011, **35**, 639–652.

References

- Alexander, C., 1999. Correlation and cointegration in energy markets. In: Kaminsky, V. (Ed.), *Managing Energy Price Risk*, second ed. Risk Publications, London, pp. 573–606.
- Arnold, L., 1974. *Stochastic Differential Equations: Theory and Applications*. John Wiley & Sons, New York.
- Bates, D., 1996. Jump and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies* 9, 69–107.

- Bessembinder, H., Coughenour, J., Seguin, P., Smoller, M., 1995. Mean-reversion in equilibrium asset prices: Evidence from the futures term structure. *Journal of Finance* 50, 361–375.
- Black, F., 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3, 167–180.
- Brennan, M., Crewe, N., 1997. Hedging long maturity commodity commitments with short-dated futures contracts. In: Dempster, M.A.H., Pliska, S.R. (Eds.), *Mathematics of Derivative Securities*. Cambridge University Press, New York, pp. 165–189.
- Carmona, R., Durrleman, V., 2003. Pricing and hedging spread options. *SIAM Review* 45, 627–685.
- Clewlow, L., Strickland, C., 1999. Valuing Energy Options in a One Factor Model Fitted to Forward Prices. Working Paper, University of Technology, Sydney, Australia.
- Dempster, M., Hong, S., 2000. Pricing spread options with the fast Fourier transform. In: Geman, H., Madan, D., Pliska, S.R., Vorst, T. (Eds.), *Mathematical Finance, Bachelier Congress*. Springer-Verlag, Berlin, pp. 203–220.
- Dicky, D., Fuller, W., 1979. Distribution of the estimates for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427–429.
- Durbin, J., Koopman, S., 1997. Monte Carlo maximum likelihood estimation of non-Gaussian state space models. *Biometrika* 84, 669–684.
- Engle, R., Granger, C., 1987. Cointegration and error correction: Representation, estimation and testing. *Econometrica* 55, 251–276.
- Eydeland, A., Geman, H., 1998. Pricing power derivatives. *RISK*, January.
- Fama, E., French, K., 1988. Permanent and temporary components of stock prices. *Journal of Political Economy* 96, 246–273.
- Fuller, W., 1976. *Introduction to Statistical Time Series*. John Wiley & Sons, New York.
- Geman, H., 2005a. *Commodities and Commodities Derivatives: Modelling and Pricing for Agriculturals Metals and Energy*. John Wiley & Sons, Chichester.
- Geman, H., 2005b. Energy commodity prices: Is mean reversion dead? *Journal of Alternative Investments* 8, 34–47.
- Geman, H., Nguyen, V., 2005. Soybean inventories and forward curves dynamics. *Management Science* 51, 1076–1091.
- Gibson, R., Schwartz, E., 1990. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance* 45, 959–976.
- Girma, P., Paulson, A., 1999. Risk arbitrage opportunities in petroleum futures spreads. *Journal of Futures Markets* 19, 931–955.
- Gregory, A., Hansen, B., 1996. Residual-based tests for cointegration in models with regime shifts. *Journal of Econometrics* 70, 99–126.
- Hamilton, J., 1994. *Times Series Analysis*. Princeton University Press.
- Harvey, A., 1989. *Forecasting Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge.
- Hilliard, J., 1999. Analytics underlying the metallgesellschaft hedge: Short term futures in a multiperiod environment. *Review of Finance and Accounting* 12, 195–219.
- Hilliard, J., Reis, J., 1998. Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot. *Journal of Financial and Quantitative Analysis* 33, 61–86.
- Hull, J., White, A., 1990. Pricing interest-rate derivative securities. *Review of Financial Studies* 3, 573–592.
- Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica* 59, 1551–1580.
- Kirk, E., 1995. Correlation in the energy markets. In: *Managing Energy Price Risk*. Risk Publications, London, pp. 71–78.
- Margrabe, W., 1978. The value of an option to exchange one asset for another. *Journal of Finance* 33, 177–186.

- Mbanefo, A., 1997. Co-movement term structure and the valuation of energy spread options. In: Dempster, M.A.H., Pliska, S.R. (Eds.), *Mathematics of Derivative Securities*. Cambridge University Press, Cambridge, pp. 99–102.
- Neuberger, A., 1998. Hedging long-term exposures with multiple short-term futures contracts. *Review of Financial Studies* 12, 429–459.
- Pan, J., 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3–50.
- Richter, M., Sorensen, C., 2002. Stochastic Volatility and Seasonality in Commodity Futures and Options: The Case of Soybeans. Working Paper, Copenhagen Business School.
- Routledge, B., Seppi, D., Spatt, C., 2000. Equilibrium forward curves for commodities. *Journal of Finance* 55, 1297–1338.
- Samuelson, P., 1964. Theoretical notes on trade problems. *Review of Economics and Statistics* 46, 145–154.
- Schwartz, E., 1997. The stochastic behaviour of commodity prices: Implications for valuation and hedging. *Journal of Finance* 52, 923–973.
- Schwartz, E., Smith, J., 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* 46, 893–911.
- Soronow, D., Morgan, C., 2002. Modelling Locational Spreads in Natural Gas Markets. Working Paper, Financial Engineering Associates. http://www.fea.com/resources/pdf/a_modeling_locational.pdf.
- Villar, J., Joutz, F., 2006. The Relationship Between Crude Oil and Natural Gas Prices. Working Paper, Department of Economics, The George Washington University.

Appendix A

Two Price Method of Simulating Spreads

In the risk neutral measure, Hilliard and Reis (1998) show that the futures price $H_i(t, T)$, $i = 1, 2$, in the Gibson–Schwartz model (5.1) follows

$$\frac{dH_i(t, T)}{H_i(t, T)} - \sigma_{i,1} dW_{i,1} - \frac{1 - e^{-k_i(T-t)}}{k_i} \sigma_{i,2} dW_{i,2}, \quad (A.1)$$

where T is the maturity date of the futures and $\rho_{12} dt = E[dW_{i,1} dW_{i,2}]$. Thus by integrating (1.1) the spot prices $S_{i,T}$ at maturity T are given by

$$S_{i,T} = H_i(T, T) = \exp\left(-\frac{1}{2}v_i^2 + v_i \cdot \epsilon_i\right), \quad (A.2)$$

where

$$\begin{aligned} v_i^2 &:= \int_t^T \left[\sigma_{i,1}^2 + \sigma_{i,2}^2 \left(\frac{1 - e^{-k_i(T-s)}}{k_i} \right)^2 - 2\rho_{12} \sigma_{i,1} \sigma_{i,2} \frac{1 - e^{-k_i(T-s)}}{k_i} \right] ds \\ &= \sigma_{i,2}^2 (T-t) - \frac{2\sigma_{i,1} \sigma_{i,2} \rho_{12}}{k_i} \left(T-t - \frac{1 - e^{-k_i(T-t)}}{k_i} \right) \\ &\quad + \frac{\sigma_{i,2}^2}{k_i^2} \left(T-t - \frac{2 - 2e^{-k_i(T-t)}}{k_i} + \frac{1 - e^{-k_i(T-t)}}{2k_i} \right) \end{aligned} \quad (A.3)$$

and ϵ_i is a normal random variable with mean zero and standard deviation one.

Defining the average volatility per unit time $\sigma_{av} := \sqrt{v^2 / (T-t)}$ in terms of the cumulative volatility of each, (A.2) shows that the simulation of each spot price S_T is exactly the same as simulating a Black (1976) driftless GBM model with volatility σ_{av} . Thus the options on S_T can also be priced by the Black

(1976) formula with σ_{av} as the input volatility. Moreover, σ_{av} can be *observed* directly from the market as the implied volatility of an option on a future with the same maturity as the futures contract. However in order to simulate the spread, we also need to know the estimated correlation $\hat{\rho}_{12}$ between ϵ_1 and ϵ_{12} which, as discussed in the introduction, is notoriously volatile and difficult to obtain. In [Sections 5.5](#) and [5.6](#) of this paper we assume $\hat{\rho}_{12}$ to be a *constant* represented by the historical correlation between the one month futures prices.

6

Analysing the Dynamics of the Refining Margin: Implications for Valuation and Hedging

Andrés García
Mirantes

Javier Población

Gregorio Serna

6.1	Introduction	96
6.2	Data and Preliminary Findings.....	98
	Data Description • Preliminary Findings about the Refining Margin	
6.3	Common Long-Term Trend Factor Models	101
	Theoretical Models • Estimation Results	
6.4	Crack-Spread Option Valuation	110
	Data • Option Valuation Methodology • Commodity Price Dynamics • Results	
6.5	Conclusion	115
	References.....	115
	Appendix A: Estimation Methodology	116
	Appendix B: Modelling the Crack-Spread Directly.....	118

It is well known that the prices of crude oil and its primary refined products (i.e. heating oil and gasoline) are cointegrated. In this paper, we extend this empirical evidence by showing that the refining margin is stationary and therefore, exhibits different dynamics from crude oil or its primary refined products. Furthermore, we show that crude oil, heating oil and gasoline are not only cointegrated but also share common long-term dynamics. This finding has crucial implications in terms of managing and hedging the risk faced by refining companies because the common long-term trend finding implies that the refining margin risk only reflects short-term effects. Specifically, in this paper, a way to hedge the refining margin with crack-spread options is analysed, and we find that assuming a common long-term trend for crude oil, heating oil and gasoline is the most accurate approach to hedging.

Keywords: Commodity markets; Commodity prices; Energy derivatives; Energy markets

6.1 Introduction

The refining process converts 47% of crude oil barrels into gasoline, 24% into diesel fuel and heating oil, 13% into jet fuel oil, 4% into heavy fuel oil, 4% into liquefied petroleum gas (LPG) and 8% into other products such as asphalt.* Like crude oil, each of these products has a market price that is quoted on organized markets. Therefore, there is a relationship between refined-product prices and crude oil prices that is known as the 'refining margin'. However, most of the articles on the stochastic behaviour of commodity prices are focused either on crude oil or on the refined products derived from crude oil rather than the refining margin itself.

In this paper, we focus on the refining margin. Contrary to oil prospecting and production, refining is a margin business; a refining company buys oil and sells refined products (e.g., gasoline and heating oil) and makes a profit that is, in principle, independent of oil and product prices. In other words, the profit of a refining company (i.e. the refining margin) is related to the difference between the prices of crude oil and the refined products. However, this difference does not necessarily decrease (or increase) when the prices of crude oil and its products rise (or fall). As a result, refineries face a significant risk when, for example, oil prices rise but product prices remain static or decline. It is clear that refining companies must somehow protect their refining margin using derivative contracts.

This risk faced by refining companies may result in unpleasant consequences that have been thoroughly studied in the literature, such as the costs of financial distress (Myers and Smith 1982, Smith and Slutz 1985), external financing costs and investment opportunities (Froot *et al.* 1993) and liquidity constraints (Holmström and Tirole 2000, Mello and Parsons 2000). However, it is important to consider the tax incentives (see, for example, Smith and Slutz [1985] and Graham and Smith [1999]). Therefore, the refining margin deserves to be independently analysed, which is the main objective of this paper.

Previous studies, including those by Serletis (1992, 1994), Pindyck (1999), Gjolberg and Johnsen (1999), Asche *et al.* (2003) and Lanza *et al.* (2005), provide evidence of unit roots and cointegration in the prices of crude oil and refined products. In this paper, this empirical evidence is extended by showing that the refining margin is stationary and, therefore, exhibits dynamics that are different from those of crude oil or its main refined products. Furthermore, we show that crude oil, heating oil and gasoline are not only cointegrated but also exhibit common long-term dynamics, implying that the refining margin reflects only short-term effects.

The hypothesis of a common long-term trend is confirmed using different factor models to jointly explain the dynamics of commodity prices. We find that the most suitable model in terms of simplicity and fit is the one that assumes a common long-term trend for crude oil, heating oil and gasoline.

This common long-term trend model for crude oil and its main refined products is framed within the family of multi-factor models proposed by Schwartz (1997) and a related series of papers, including those by Schwartz and Smith (2000), Cortazar and Schwartz (2003) and Cortazar and Naranjo (2006). All of these multi-factor models assume that the spot price is the sum of both short- and long-term components. Long-term factors account for the long-term dynamics of commodity prices, which are assumed to follow a random walk, whereas the short-term factors account for the mean-reverting components in commodity prices. Moreover, in the cases of heating oil and gasoline, a deterministic seasonal component is added, as suggested by Sørensen (2002).†

Cortazar *et al.* (2008) have already proposed and estimated a model in which several commodities can exhibit both common and specific factors. They presented their multi-commodity model as a way to improve the estimation of futures prices for a commodity with scarce data using long-maturity futures

* See the Oil Market Report (2006) developed by the International Energy Agency for more information about these issues.

† In contrast to crude oil, heating oil and gasoline exhibit a strong seasonal behaviour (see, for example, Garcia *et al.* [2012]). Therefore, the model for heating oil and gasoline must account for this fact.

prices that are available for another commodity. They applied their model to WTI-Brent* crude oil and WTI-Unleaded gasoline futures contracts but not to heating oil.

However, in contrast to the study by Cortazar *et al.* (2008), the primary objective of this paper is to model and hedge the refining margin. Specifically, we show that the refining margin exhibits dynamics that are different from those of crude oil and its main refined products because the refining margin only reflects short-term effects. Furthermore, we show that the proper way to model and hedge the refining margin is to consider the common long-term factor model because it provides a more accurate crack-spread option valuation; specifically, crack-spread option quotes imply that the refining margin is stationary, and therefore, commodities exhibit a common long-term trend. Consequently, the refining margin only reflects short-term effects; thus, a multi-commodity model with a common long-term trend is the best way to model and hedge it. In this sense, Cortazar *et al.* (2008) suggested that if the spreads between two commodities are of interest, the use of multi-commodity models should provide much more stable estimates, particularly for the commodity for which data are scarce.

The fact that crude oil, heating oil and gasoline share a common long-term trend will have straightforward implications for managing and hedging refining margin risk, specifically in terms of crack-spread option valuation. Crack-spread options are used to protect the refining margin while simultaneously allowing market participants to take advantage of favourable changes in the spread. A crack-spread call option is a contract that gives the holder the right (but not the obligation) to buy a refined-product futures contract from the writer and sell the writer a crude oil futures contract, paying a previously agreed upon crack-spread price. A crack-spread put allows one to sell a refined product futures contract and buy a crude oil futures contract, paying a crack-spread price.

The primary organized market for crack-spread options is the NYMEX, which is now included in CME Group, where calls and puts at a one-to-one ratio between the New York Harbour unleaded gasoline or heating oil futures contract and the light sweet crude oil futures contract can be found.

Given that these crack-spread options involve several assets, the techniques used in their valuation are complex. Dempster *et al.* (2008) propose a model for valuing spread options on two commodity prices that are cointegrated and give two applications to a set of theoretical (not traded) European options on the crack-spread between heating oil and WTI crude oil and the location spread between Brent blend and WTI crude oil. Dempster *et al.* (2008) point out that the correlation between two asset returns is hard to model (Kirk 1995, Mbanefo 1997, Alexander 1999), and therefore, the crack-spread should be modelled directly. However, from our point of view, there are other reasons that must be taken into account when modelling the crack-spread. For example, as suggested by Cortazar *et al.* (2008), using two or more price series, it is possible to extract more information for estimation purposes. Therefore, a common long-term trend model for the two assets involved in the crack-spread can be an alternative to the approach suggested by Dempster *et al.* (2008).

In this paper, we present a simplified model for valuing this type of option that assumes a common long-term trend for the prices of crude oil and its main refined products (i.e. heating oil and gasoline). As opposed to Dempster *et al.* (2008), who use a theoretical (not traded) option database, we use an extensive database of American crack-spread options (heating oil vs. WTI and unleaded gasoline vs. WTI) traded at NYMEX to find that by assuming a common long-term trend for the prices of crude oil and refined products, more accurate option valuations can be obtained than when using models with more factors and parameters, including the one proposed by Dempster *et al.* (2008). This result also implies that the advantage of the model with a common long-term trend, i.e. by using two or more price series, it is possible to extract more information for estimation purposes, has more importance in terms of

* WTI (West Texas Intermediate), also known as Texas light sweet, is a type of crude oil used as a benchmark in oil pricing and is the underlying commodity of New York Mercantile Exchange's oil futures contracts. Brent is another type of crude oil, which is sourced from the North Sea.

valuation errors than the advantage of the Dempster *et al.* (2008) model, i.e. avoiding the need to model the correlation between two asset returns.

Even more importantly, when assuming a common long-term trend, theoretical option values are lower (and closer to the real values) than when one assumes different long-term trends for each commodity. This result may be related to the fact that the option holder assumes lower risk with a common long-term trend than with different long-term trends for each commodity because, with common long-term trends, the refining margin risk reflects only short-term effects. This result can be explained as follows. Given that the refining margin is proved to be stationary, its behaviour can be better explained with a model without long-term effects because long-term effects are not stationary. By using a common long-term trend model, the (common) long-term factors disappear when computing the refining margin, which does not occur when we use the model with different long-term trends. Therefore, because the refining margin only reflects short-term effects, its volatility is bounded in time.* Consequently, as investors face lower risk, option prices must be lower.

Therefore, we can conclude that the preferred model for valuing this type of option is one that assumes a common long-term trend. Additionally, the evidence of a common long-term trend is useful not only in the valuation and hedging of commodity-contingent claims but also in defining procedures for evaluating investment projects related to natural resources, particularly when determining optimal investment rules.

The remainder of this paper is organized as follows. Section 6.2 presents the data and some preliminary findings regarding the refining margin. We show that in contrast to crude oil and its main refined products, the refining margin is stationary, implying that it exhibits dynamics that are different from crude oil and its products. In Section 6.3, we show that crude oil, heating oil and gasoline are not only cointegrated but also exhibit common long-term trends, implying that the refining margin only reflects short-term effects. In Section 6.4, we use the results from Sections 6.2 and 6.3 to value crack-spread options, and we analyse the implications for hedging the refining margin. Specifically, Section 6.4 shows the valuation results of the crack-spread options listed on NYMEX using a model that assumes a common long-term trend for crude oil and the main refined product prices, a model that allows for a long-term trend for each commodity, a model that postulates uncorrelated sub-models for each commodity and the model by Dempster *et al.* (2008). Finally, Section 6.5 concludes with a summary and discussion.

6.2 Data and Preliminary Findings

In this section, we present some preliminary findings regarding the refining margin. Specifically, we show that the refining margin exhibits dynamics that are different from those of crude oil and its main refined products, implying that refining companies must use specific hedging products whenever they choose to cap their losses on the refining margin.

However, before presenting these preliminary findings regarding the refining margin, it is useful to briefly describe the data that will be used in this and the following sections.

6.2.1 Data Description

The data set used in this paper consists of weekly observations of WTI (light sweet) crude oil, heating oil and unleaded gasoline futures prices traded at NYMEX[†] and weekly observations of the refining margin for WTI oil on the U.S. Gulf coast.

* If the refining margin were to reflect long-term effects, which are not stationary, its volatility would tend to increase with time.

[†] Details about the contracts can be found on the NYMEX homepage.

WTI crude oil futures with maturities of 1 month to 7 years, heating oil futures with maturities from 1 to 18 months and gasoline futures with maturities of 1 to 12 months are currently being traded at NYMEX. However, in the case of gasoline, there is not enough liquidity for the futures with longer maturities. Therefore, in the cases of WTI crude oil and heating oil, our data set comprises futures prices from 1 to 18 months (664 weekly observations), from 9/9/1996 to 5/25/2009, and in the case of RBOB* gasoline, our data set comprises futures prices from 1 to 9 months (622 weekly observations), from 6/30/1997 to 5/25/2009.

With regard to the refining margin data, it should be noted that we can use different refining margin series depending on the type of crude oil used (WTI, Brent, Dubai, etc.) and the geographical area represented (U.S. Gulf Coast, Rotterdam, Singapore, etc.). However, to cohere with the NYMEX crude oil and refined products data described above, we will use the refining margin for WTI crude oil on the U.S. Gulf coast.[†]

The refining margin series used in this paper was calculated (as suggested by the International Energy Agency[‡]) as the gross product worth of a refinery on the U.S. Gulf Coast using the catalytic cracking method, minus the costs of crude oil and freight. The gross product worth is the weighted average value of all refined product components of a barrel of crude oil, computed by multiplying the spot price of each product by its percentage share of the yield of the barrel. Weekly observations of this refining margin from 1/26/1997 to 5/24/2009 (representing 643 observations) were used in the tests. The main descriptive statistics of the series are summarized in Table 6.1.

TABLE 6.1 Descriptive Statistics

	WTI Crude Oil		Heating Oil		Gasoline		Ref. Margin	
	Mean	Volatility (%)	Mean	Volatility (%)	Mean	Volatility (%)	Mean	Stand. Dev.
Spot							3.58	5.21
F1	41.76	32.14	48.94	31.51	51.16	36.99		
F2	41.89	29.32	49.11	29.15	50.99	31.99		
F3	41.92	27.56	49.21	27.49	50.82	29.60		
F4	41.87	26.09	49.23	26.22	50.61	26.59		
F5	41.80	24.79	49.20	25.03	50.36	25.28		
F6	41.71	23.71	49.14	23.89	50.18	24.36		
F7	41.60	22.83	49.08	22.73	50.03	23.57		
F8	41.50	22.09	49.00	21.76	49.93	23.60		
F9	41.40	21.45	48.91	20.81	49.85	23.22		
F10	41.30	20.87	48.81	20.08				
F11	41.20	20.35	48.71	19.58				
F12	41.11	19.88	48.60	19.12				
F13	41.02	19.45	48.52	18.83				
F14	40.93	19.06	48.45	18.62				
F15	40.84	18.72	48.39	18.53				
F16	40.76	18.43	48.33	18.59				
F17	40.69	18.16	48.28	18.56				
F18	40.62	17.91	48.24	18.48				

The table shows the mean and volatility of the four commodity series prices. F1 is the futures contract closest to maturity, F2 is the contract second-closest to maturity and so on.

* RBOB stands for 'reformulated gasoline blendstock for oxygen blending'. It is the benchmark gasoline contract on the New York Mercantile Exchange.

[†] The results do not change if another refining margin is used. For the sake of brevity, the results for other refining margins are not presented here, although they are available from the authors upon request.

[‡] Oil Market Report, Annual Statistical Supplement. International Energy Agency. <http://www.omrpublic.iea.org>; 2006.

6.2.2 Preliminary Findings about the Refining Margin

Previous studies, including those by Serletis (1992, 1994), Pindyck (1999), Gjolberg and Johnsen (1999), Asche *et al.* (2003) and Lanza *et al.* (2005), provide evidence of unit roots and cointegration in the prices of crude oil and refined products (i.e. heating oil and gasoline). However, to the best of our knowledge, there is no evidence of stationarity of the refining margin, which is composed not only of crude oil, heating oil and gasoline but also of other products such as liquefied petroleum gas (LPG), asphalt and kerosene.

Table 6.2 summarizes the results of the unit root tests for WTI crude oil, heating oil and gasoline prices: and the refining margin. The empirical evidence from previous studies of a unit root in crude oil, heating oil and gasoline prices is confirmed in the present work using the standard Augmented Dickey-Fuller and Phillips-Perron tests. However, the refining margin does not show evidence of a unit root.*† Therefore, the refining margin seems to show different dynamics than that of crude oil and its main refined products.

This result has important implications in terms of hedging the position of a refining company. Refining companies face risk associated with a refining margin that is different from that of crude oil and its main refined products. Thus, refining companies must use specific hedging products that are different from those used to hedge the risk associated with crude oil and its main refined products. As discussed in the introduction, crack-spread options meet this need.

Moreover, in the next section, we show that crude oil, heating oil and gasoline prices are not only cointegrated but also exhibit common long-term dynamics.‡ This finding has important implications in terms of crack-spread option valuation, as will be discussed in [Section 6.4](#).

TABLE 6.2 Unit Root Tests.

Series	ADF	Phillips–Perron
WTI crude oil	–1.2409 (0.6582)	–1.6380 (0.4625)
Heating oil	–1.8821 (0.3409)	–1.5668 (0.4991)
Gasoline	–1.8517 (0.3554)	–1.8534 (0.3546)
Refining margin	–5.9062*** (0.0000)	–7.7892*** (0.0000)

The table shows the statistics of the Augmented Dickey–Fuller (ADF) and Phillips–Perron unit root tests. MacKinnon *p*-values are shown in parentheses. The results are reported with * indicating rejection of the null hypothesis of a unit root at a 10% significance level, ** indicating rejection at 5% and *** at indicating rejection at 1%.

* To the best of the authors' knowledge, there are no spot prices for gasoline, heating oil or crude oil associated with the futures traded at NYMEX. Therefore, weekly observations of 1-month futures prices for WTI crude oil, heating oil and gasoline are used as proxies for spot prices.

† As discussed below, heating oil and gasoline show a seasonal effect, whereas crude oil does not (Garcia *et al.* 2012). Therefore, the refining margin must inherit a seasonal component as well. Thus, we have repeated the unit root tests after removing the seasonal component in the series, obtaining very similar results to those presented in Table 6.2.

‡ For brevity's sake, the results of the cointegration tests are not presented here, although they are available from the authors upon request.

6.3 Common Long-Term Trend Factor Models

In this section, we propose and estimate different factor models either with or without assuming common long-term trends for crude oil, heating oil and gasoline. These comparisons will demonstrate that the most suitable model in terms of both simplicity and fit is the one that assumes a common long-term trend for all three commodities. This result suggests that the three commodities are not only cointegrated, as shown in previous studies, but also share a common long-term trend, implying that the refining margin only reflects short-term effects. This finding will have straightforward implications for managing and hedging the refining margin risk, specifically in terms of crack-spread option valuation.

Given that the main objective of this paper is to characterize the dynamics of the refining margin, the possibility of modelling the refining margin directly instead of modelling each price series as a stochastic system has been considered. However, by modelling each commodity as a stochastic system, we use richer information than directly modelling the price differences because in the former case, we account for all of the information contained in the price series. In contrast, by modelling the price differences directly, we would lose some of the information contained in the original price series because one series must be subtracted from the other.*

Furthermore, by modelling each commodity as a stochastic system, we will be able to show that all three commodities share a common long-term trend, suggesting that the refining margin only reflects short-term effects.

It seems clear that modelling each commodity separately is the way to obtain the best fit for a given data set. However, if we were to obtain a similar goodness of fit when modelling the three commodities jointly with a common long-term trend, the conclusion would be that all three commodities share a common long-term trend. It is also possible to compare the results obtained from modelling the commodities jointly with and without the assumption of a common long-term trend; if there is a common long-term trend, the results must be comparable.

This result has important implications in terms of hedging the position of a refining company. Refining companies face risk associated with a refining margin that

6.3.1 Theoretical Models

Here, we present three different specifications used to model the stochastic behaviour of the three commodities under study within the context of the two-factor model proposed by Schwartz and Smith (2000). Given the existing empirical evidence,[†] this is a reasonable approach to this kind of commodity. In this model, the log-spot price[‡] (X_t) is assumed to be the sum of two stochastic factors: a short-term deviation (χ_t) and a long-term equilibrium price level (ξ_t). Thus,

$$X_t = \xi_t + \chi_t. \quad (6.1)$$

The stochastic differential equations (SDEs) for these factors are as follows:

$$\begin{cases} d\xi_t = \mu_\xi dt + \sigma_\xi dW_{\xi t}, \\ d\chi_t = -\kappa\chi_t dt + \sigma_\chi dW_{\chi t}, \end{cases} \quad (6.2)$$

* Specifically, by modelling each price series separately, it is possible to use futures prices with different maturities for each commodity. This is one of the main advantages of the Cortazar *et al.* (2008) multi-commodity model. However, modelling directly the price differences would force us to use futures with the same maturities for both commodities.

[†] See, for example, Schwartz (1997).

[‡] For the sake of simplicity, it is typical in the literature to formulate the model in terms of log-spot prices because log-spot prices are assumed to be normally distributed, whereas spot prices are assumed to be log-normally distributed.

where dW_{ξ_t} and dW_{χ_t} can be correlated ($dW_{\xi_t} dW_{\chi_t} = \rho_{\xi\chi} dt$) and $\rho_{\xi\chi}$ represents the coefficient of correlation between long- and short-term factors.

In this model, μ_{ξ} and σ_{ξ} represent the trend and volatility, respectively, of the long-term factor, whereas κ and σ_{χ} represent the speed of adjustment and volatility, respectively, of the short-term factor.

This model captures the most important features of commodity prices. Specifically, as explained by Schwartz and Smith (2000), the equilibrium price level (ξ_t) is assumed to follow a geometric Brownian motion. In this motion, the drift reflects market expectations about the exhaustion of existing supplies, improvements in technology for the production and discovery of the commodity and inflation or political/regulatory effects. The short-term deviations (χ_t) are assumed to revert towards zero following an Ornstein–Uhlenbeck process. These deviations reflect short-term changes in demand, which are attenuated by market participants adjusting their inventories, accounting for the commonly observed mean-reverting feature of commodity prices.

Moreover, in the cases of heating oil and gasoline, a deterministic seasonal component is added, as suggested by Sorensen (2002).^{*} Therefore, the log spot price for heating oil and gasoline (X_t) is assumed to be the sum of two stochastic factors (χ_t and ξ_t) and a deterministic seasonal trigonometric component (α_t) (i.e. $X_t = \xi_t + \chi_t + \alpha_t$). The SDEs for ξ_t and χ_t are given by Equation (6.2) and by

$$d\alpha_t = 2\pi\phi\alpha_t^* dt \quad \text{and} \quad d\alpha_t^* = -2\pi\phi\alpha_t dt,$$

where α_t^* is the other seasonal factor that complements α_t , and ϕ is the seasonal period.

The first specification, in which each of the three commodities is modelled separately, will be the simplest one; it is represented by a six-factor model with no correlation between the factors. However, this model is not very realistic because the no-correlation assumption is clearly an undesirable property in valuing commodity-contingent claims, as discussed below.

To solve this problem, we propose a second six-factor model, a joint model for all three commodities that allows for correlation between factors. In this model, the log-spot price of crude oil (X_{1t}) is assumed to be the sum of two stochastic factors: a short-term deviation (χ_{1t}) and a long-term equilibrium price level (ξ_{1t}) (i.e. $X_{1t} = \xi_{1t} + \chi_{1t}$). However, in the cases of heating oil and gasoline, a deterministic seasonal component, α_{it} , $i = 2, 3$ is added, where subscripts 2 and 3 refer to heating oil and gasoline, respectively. Therefore, the log-spot price for heating oil and gasoline (X_{it}) will be $X_{it} = \xi_{it} + \chi_{it} + \alpha_{it}$, $i = 2, 3$. The SDEs of the factors for this joint model without a common long-term trend are

$$\left. \begin{aligned} d\xi_{it} &= \mu_{\xi_i} dt + \sigma_{\xi_i} dW_{\xi_{it}}, \\ d\chi_{it} &= -\kappa_i \chi_{it} dt + \delta_{\chi_i} dW_{\chi_{it}}, \end{aligned} \right\} \quad i=1, 2, 3$$

$$\left. \begin{aligned} d\alpha_{it} &= 2\pi\phi_i \alpha_{it}^* dt, \\ d\alpha_{it}^* &= -2\pi\phi_i \alpha_{it} dt, \end{aligned} \right\} \quad i=2, 3,$$
(6.3)

where $dW_{\xi_{1t}}$, $dW_{\xi_{2t}}$, $dW_{\xi_{3t}}$, $dW_{\chi_{1t}}$, $dW_{\chi_{2t}}$ and $dW_{\chi_{3t}}$ can show any correlation structure, resulting in 15 correlation parameters.

This second model does account for the relationships between series, but it does so in a somewhat ambiguous way. We have 15 correlations to consider, none of which are negligible. We have three correlated long-term trends, though the relationship between series does not stop there. We cannot take this correlation as the only measure because the long-term trend for crude oil is also correlated with the short-term trend for refined products. Moreover, questions regarding subjects like the general market

^{*} Sorensen (2002) suggested introducing a deterministic seasonal component into models for agricultural commodities. Here, we use Sorensen's proposal for heating oil and gasoline, which exhibits strong seasonal behaviour (see, for example Garcia *et al.* 2012).

trend cannot be answered. Therefore, in our second model, such questions are meaningless unless we assume that some combination of long-term trends is 'representative'.

This problem can be solved by means of a third specification with only one long-term trend for all three commodities. Using such a model, questions such as the one described above can be fully answered: the general trend is the common long-term trend. Furthermore, it is even possible to see the relationship between the general trend and each of the series by observing its long-run/short-run correlation coefficient. In addition, the (isolated) influence of one series on another can be directly determined by observing the short-term/short-term correlation coefficient. In this model, the log-spot price of crude oil (X_{1t}) is assumed to be the sum of two stochastic factors: a short-term deviation (χ_{1t}), which is different for each commodity; and a common long-term equilibrium price level (ξ_t), where $X_{1t} = \xi_t + \chi_{1t}$. However, in the cases of heating oil and gasoline, a deterministic seasonal component, α_{it} , $i = 2, 3$, is also added. Therefore, the log-spot price for heating oil and gasoline (X_{it}) will be $X_{it} = \xi_t + \chi_{it} + \alpha_{it}$, $i = 2, 3$. The SDEs of the factors for this joint model with a common long-term trend are

$$\begin{aligned} d\xi_t &= \mu_\xi dt + \sigma_\xi dW_{\xi_t}, \\ d\chi_{it} &= -\kappa_i \chi_{it} dt + \sigma_{\chi_i} dW_{\chi_{it}}, \quad i=1, 2, 3, \\ d\alpha_{it} &= 2\pi\phi\alpha_{it}^* dt, \\ d\alpha_{it}^* &= -2\pi\phi\alpha_{it} dt, \end{aligned} \quad \left. \vphantom{\begin{aligned} d\chi_{it} \\ d\alpha_{it} \\ d\alpha_{it}^* \end{aligned}} \right\} \quad i=2, 3, \quad (6.4)$$

where dW_{ξ_t} , $dW_{\chi_{1t}}$, $dW_{\chi_{2t}}$ and $dW_{\chi_{3t}}$ can show any correlation structures resulting in six correlation parameters.

This common long-term trend model requires two additional parameters to account for the variable quality of refining products. Therefore, even if their long-term dynamics are the same, their price levels and the effects of the long-term factor on their prices may differ. Specifically, because of differences in quality, the final price level should be different for crude oil than for refined products as well as for different refined products, if one is compared to the other. Assuming a common long-term trend for crude oil and refining products, these price level differences must be stationary; thus, they have been included in the short-term component, which is different from one commodity to the next. However, because this short-term factor has no drift and affects the log-spot price rather than the price level, it is necessary to introduce a constant (K_i) in the price level to account for this fact. Furthermore, these quality differences might lead to differences in the way that this common long-term trend affects the price dynamics of each commodity. Thus, because the long-term factor is the same for the three commodities, this factor will be multiplied by a different constant (C_i) for each commodity.

Therefore, the spot price for crude oil can be calculated as $P_{1t} = \exp(\xi_t + \chi_{1t})$.^{*} However, in the cases of heating oil and gasoline, the spot price will be $P_{it} = K_i + \exp(C_i \cdot \xi_t + \chi_{it} + \alpha_{it})$, $i = 2, 3$.

The third model (i.e. the joint model with a common long-term trend) is preferable to the second one (i.e. the joint model without a common long-term trend); it contains fewer parameters and is therefore simpler, and it has only one long-term factor, which is an advantage in valuing long-term commodity-contingent claims.

As stated above, it is clear that the first specification, in which each commodity is modelled separately, must be the model that best fits the data. Nonetheless, if we achieve a similar goodness of fit when modelling the three commodities jointly using the third specification, the conclusion should be that all three commodities share a common long-term trend. If this is the case, the preferable specification will be the third one because it is the simplest model. Moreover, the joint model with a common long-term trend is the most appropriate model for valuing options involving several commodity prices, as we will discuss in [Section 6.4](#).

^{*} To avoid identification problems, in the case of crude oil, both constants are assumed to be equal to zero. Thus, the other commodity constants must be interpreted in terms of differences from crude oil.

It is also possible to test the three models proposed above with the commodities in pairs (i.e. crude oil and heating oil, crude oil and gasoline or heating oil and gasoline). These specifications in pairs will be useful for valuing the crack-spread options in [Section 6.4](#).

Finally, it should be noted that the differences between the models presented above lie in the number of factors and the correlation assumed between them.

6.3.2 Estimation Results

Here, we present the results of the estimation of the three models with and without the assumption of a common long-term trend. The data set used in this section consists of weekly observations of WTI (light sweet) crude oil, heating oil and unleaded gasoline futures prices traded at NYMEX, as described in [Section 6.2.1](#). When using models with more than one commodity, we have chosen to maintain a consistent time to maturity between futures contracts to avoid decompensating the short-term/long-term relations.* Therefore, different data sets have been used to estimate the parameters for the different models presented above.

Different data sets for gasoline and the two other commodities were used because of liquidity constraints. In the case of gasoline, as stated above, the available futures contracts are less liquid, and their maturities are shorter than those of the other two commodities' contracts. Consequently, given that in the case of heating oil there exists liquidity to support futures contracts with up to 18 months to maturity, we decided to use a data set with more futures contracts and with futures contracts that have longer maturities than gasoline (for which there is not enough liquidity for futures with maturities longer than 9 months). In the case of crude oil, futures contracts with up to 7 years to maturity are available. Nevertheless, we decided to use crude oil futures contracts with the same maturity as the ones used for the other commodities to avoid decompensating the short-term/long-term effects. Schwartz (1997) realized that mean-reversion effects tend to be smaller for contracts with longer maturities. García *et al.* (2010) found evidence suggesting the same conclusion in the case of natural gas. Thus, to avoid undesirable effects, we made the time to maturity for futures contracts consistent across all commodities in models that consider more than one commodity.

Specifically, in the case of RBOB gasoline, the data set comprises contracts F1, F3, F5, F7 and F9 from 06/30/1997 to 05/25/2009, yielding 622 quotations for each contract. In this case, for instance, F1 is the contract for the month closest to maturity and F2 is the contract for the second-closest month to maturity.

There are two data sets for WTI crude oil and heating oil. The first data set comprises contracts F1, F3, F5, F7 and F9 from 06/30/1997 to 05/25/2009, yielding 622 quotations for each contract. The second data set comprises F1, F4, F7, F11, F15 and F18 from 09/09/1996 to 05/25/2009, yielding 664 quotations for each contract. The data set we employed depended on the specific case: the first data set was used when either crude oil or heating oil was used jointly with gasoline; and the second data set was used in all other situations. As explained in Schwartz (1997), because futures contracts have a fixed maturity date, the time to maturity changes as time progresses, though it remains in a narrow time interval. Therefore, as in Schwartz (1997), it is assumed that time to maturity does not change over time and is equal to, for instance, 1 month for F1 and 2 months for F2.

The models presented in [Section 6.3.1](#) were estimated using the Kalman filter methodology, which is briefly described in [Appendix A](#).

[Table 6.3](#) presents the results for the first specification (i.e. the two-factor model by Schwartz and Smith [2000]) applied to each commodity (i.e. WTI crude oil, heating oil and gasoline) separately and using the different data sets described above. [Table 6.4](#) presents the results of the second specification (i.e. the joint model without a common long-term trend). The results of the estimation of the third specification (i.e. the joint model with a common long-term trend for all three commodities) are presented in

* Similar results have been obtained with futures with different maturities.

TABLE 6.3 The Two-Factor Model for Each Commodity Separately

	WTI Crude Oil	WTI Crude Oil	Heating Oil	Heating Oil	Gasoline
Contracts	F1, F3, F5, F7 and F9	F1, F4, F7, F11, F15 and F18	F1, F3, F5, F7 and F9	F1, F4, F7, F11, F15 and F18	F1, F3, F5, F7 and F9
Period	06/30/1997 to 05/25/2009	09/09/1996 to 05/25/2009	06/30/1997 to 05/25/2009	09/09/1996 to 05/25/2009	06/30/1997 to 05/25/2009
Number obs	622	664	622	664	622
μ_{ξ}	0.1031 ** (0.0417)	0.1110*** (0.0339)	0.0998*** (0.0329)	0.1093*** (0.0322)	0.1830*** (0.0362)
κ	1.6738 *** (0.0266)	1.1599 *** (0.0110)	1.4484*** (0.0417)	1.5870*** (0.0185)	1.7403*** (0.0604)
σ_{ξ}	0.2118*** (0.0044)	0.1792*** (0.0036)	0.2204*** (0.0053)	0.1758*** (0.0036)	0.2245*** (0.0053)
σ_{χ}	0.2951 *** (0.0068)	0.3022*** (0.0067)	0.3152*** (0.0084)	0.2830*** (0.0067)	0.3919*** (0.0104)
$\rho_{\xi\chi}$	-0.0503 (0.0324)	-0.0360 (0.0306)	-0.0783* (0.0401)	0.1556*** (0.0313)	-0.2566*** (0.0411)
λ_{ξ}	0.2134*** (0.0418)	0.1678*** (0.0339)	0.0929*** (0.0333)	0.1341*** (0.0322)	0.2474*** (0.0373)
λ_{χ}	-0.1010* (0.0572)	-0.0245 (0.0568)	-0.4998*** (0.0491)	0.1991*** (0.0502)	0.2412*** (0.0638)
Φ			1.0042*** (0.0003)	0.9954*** (0.0002)	1.0002*** (0.0003)
σ_{η}	0.0071*** (0.0001)	0.0095*** (0.0001)	0.0115*** (0.0001)	0.0177*** (0.0002)	0.0161*** (0.0002)
Log-likelihood	23556.22	29212.38	21543.49	25668.93	19913.68
AIC	23540.22	29196.38	21525.49	25650.93	19895.68
SIC	23504.76	29160.39	21485.59	25610.44	19855.78

The table presents the results for the Schwartz and Smith (2000) two-factor model for each commodity separately (first specification). Standard errors are in parentheses. The estimated values are reported with * denoting significance at 10%, ** denoting significance at 5% and *** denoting significance at 1%.

TABLE 6.4 The Joint Model Without a Common Long-Term Trend for the Three Commodities

WTI Crude Oil (WTI), Heating Oil (HO) and Unleaded Gasoline (UG) Contracts F1, F3, F5, F7 and F9 Period 06/30/1997 to 05/25/2009 Number Obs. 622

$\mu_{\xi\text{WTI}}$	0.1056*** (0.0257)	Φ_{HO}	1.0021 *** (0.0005)
$\mu_{\xi\text{HO}}$	0.1049** (0.0418)	Φ_{UG}	1.0003*** (0.0003)
$\mu_{\xi\text{UG}}$	0.1601*** (0.0231)	$\rho_{\xi\text{WTI}\xi\text{HO}}$	0.7417*** (0.0032)
κ_{WTI}	2.3492*** (0.0495)	$\rho_{\xi\text{WTI}\xi\text{UG}}$	0.8620*** (0.0054)
κ_{HO}	1.4107*** (0.0789)	$\rho_{\xi\text{WTI}\xi\text{WTI}}$	0.1133*** (0.0074)
κ_{UG}	2.403*** (0.0500)	$\rho_{\xi\text{WTI}\xi\text{HO}}$	0.3901*** (0.0072)
$\sigma_{\xi\text{WTI}}$	0.2372*** (0.0018)	$\rho_{\xi\text{WTI}\xi\text{UG}}$	0.3353*** (0.0081)
$\sigma_{\xi\text{HO}}$	0.2200*** (0.0032)	$\rho_{\xi\text{HO}\xi\text{UG}}$	0.6624*** (0.0055)
$\sigma_{\xi\text{UG}}$	0.2275*** (0.0050)	$\rho_{\xi\text{HO}\xi\text{WTI}}$	0.1184*** (0.0071)
$\sigma_{\chi\text{WTI}}$	0.2754*** (0.0068)	$\rho_{\xi\text{HO}\xi\text{HO}}$	-0.1665*** (0.0144)
$\sigma_{\chi\text{HO}}$	0.3637*** (0.0142)	$\rho_{\xi\text{HO}\xi\text{UG}}$	0.1663*** (0.0107)
$\sigma_{\chi\text{UG}}$	0.3891*** (0.0100)	$\rho_{\xi\text{UG}\xi\text{WTI}}$	0.2513*** (0.0069)
$\lambda_{\xi\text{WTI}}$	0.0841*** (0.0258)	$\rho_{\xi\text{UG}\xi\text{HO}}$	0.3563*** (0.0075)
$\lambda_{\xi\text{HO}}$	0.0775* (0.0441)	$\rho_{\xi\text{UG}\xi\text{UG}}$	0.0943*** (0.0091)
$\lambda_{\xi\text{UG}}$	0.1398*** (0.0248)	$\rho_{\chi\text{WTI}\xi\text{HO}}$	0.4240*** (0.0092)
$\lambda_{\chi\text{WTI}}$	0.2070*** (0.0657)	$\rho_{\chi\text{WTI}\xi\text{UG}}$	0.4679*** (0.0099)
$\lambda_{\chi\text{HO}}$	-0.1381* (0.0774)	$\rho_{\chi\text{HO}\xi\text{UG}}$	0.4957*** (0.0187)
$\lambda_{\chi\text{UG}}$	0.0977** (0.0495)	σ_{η}	0.0121*** (0.0001)
Log-likelihood		63042.16	
AIC		62970.16	
SIC		62810.57	

The table presents the results obtained using the Schwartz and Smith (2000) two-factor model without a common long-term trend for the three commodities (second specification). Standard errors are in parentheses. The estimated values are reported with * denoting significance at 10%, ** denoting significance at 5% and *** denoting significance at 1%.

Table 6.5. Finally, Tables 6.6 and 6.7 present the results obtained when using the second and third specifications, respectively, for pairs of commodities (as will be described in Section 6.4). As stated above, in models with more than one commodity, we have chosen to use futures contracts with the same maturities for each commodity.

The first notable observation is that both the long-term trend (μ_ξ) and the speed of adjustment (κ) are positive and significantly different from zero in all cases, implying long-term growth and mean reversion in the commodity prices. This observation is consistent with the results obtained by Schwartz (1997) in the case of oil. Moreover, for each commodity, the estimated values of the long-term mean and volatility (μ_{ξ_i} and σ_{ξ_i} , respectively) are very similar for all specifications.

Gasoline prices exhibit a higher long-term mean and are more volatile than are the other two commodity prices; in Tables 6.3 through 6.7, the volatility coefficients are higher when gasoline appears in the model. This phenomenon can also be observed in Table 6.1. It is also interesting to note that in all cases, short-term volatility (σ_χ) is higher than long-term volatility (σ_ξ). This result is consistent with the results obtained by Schwartz and Smith (2000) for oil and García *et al.* (2012) for heating oil and gasoline.

In general, the market prices of risk associated with the long- and short-term factors (λ_ξ and λ_χ , respectively) are significantly different from zero, suggesting that the risk associated with these factors cannot be diversified. Additionally, the values of the market prices of risk are higher when the factor models are estimated separately (Table 6.3) than when the model is estimated jointly, with or without a common long-term trend (Tables 6.4 through 6.7). These results suggest that the risk associated with the long- and short-term factors is more difficult to diversify when we ignore the relationships among factors.

As expected, in the case of heating oil and gasoline, the seasonal period (φ) is roughly 1 year in all cases; this result is consistent with the findings of García *et al.* (2012).

In relation to the purpose of this paper, two important results must be highlighted. First, in the specifications without a common long-term trend (Tables 6.4 and 6.6), the values of the coefficients of correlation between long-term trends ($\rho_{\xi_i\xi_j}$) are quite large, suggesting a strong relationship between long-term trends.

TABLE 6.5 The Joint Model With a Common Long-Term Trend for the Three Commodities

WTI Crude Oil (WTI), Heating Oil (HO) and Unleaded Gasoline (UG) Contracts F1, F3, F5, F7 and F9 Period 06/30/1997 to 05/25/2009 Number Obs. 622

μ_ξ	0.1236** (0.0502)	φ_{HO}	0.9945*** (0.0007)
κ_{WTI}	1.1960*** (0.0667)	φ_{UG}	1.0055*** (0.0003)
κ_{HO}	1.2488*** (0.0666)	$\rho_{\xi\chi WTI}$	-0.2168*** (0.0102)
κ_{UG}	1.0558*** (0.0461)	$\rho_{\xi\chi HO}$	-0.2385*** (0.0072)
σ_ξ	0.2028*** (0.0043)	$\rho_{\xi\chi UG}$	-0.2583*** (0.0168)
$\sigma_{\chi WTI}$	0.3606*** (0.0054)	$\rho_{\chi WTI\chi HO}$	0.8711*** (0.0085)
$\sigma_{\chi HO}$	0.3579*** (0.0083)	$\rho_{\chi WTI\chi UG}$	0.8802*** (0.0251)
$\sigma_{\chi UG}$	0.4092*** (0.0182)	$\rho_{\chi HO\chi UG}$	0.8566*** (0.0167)
λ_ξ	0.0908* (0.0524)	κ_{HO}	1.4165** (0.5992)
$\lambda_{\chi WTI}$	0.2831** (0.1281)	κ_{UG}	1.5147*** (0.1698)
$\lambda_{\chi HO}$	-0.1011 (0.0857)	C_{HO}	0.9721 *** (0.0133)
$\lambda_{\chi UG}$	0.0084 (0.1173)	C_{UG}	0.9668*** (0.0045)
σ_η	0.0169*** (0.0001)		
Log-likelihood		62125.99	
AIC		62075.99	
SIC		61965.17	

The table presents the results for the Schwartz and Smith (2000) two-factor model assuming a common long-term trend for all three commodities (third specification). Standard errors are in parentheses. The estimated values are reported with * denoting significance at 10%, ** denoting significance at 5%, and *** denoting significance at 1%.

TABLE 6.6 The Joint Model Without a Common Long-Term Trend for Pairs of Commodities.

	WTI Crude Oil (<i>i</i>) and Heating Oil (<i>j</i>)	WTI Crude Oil (<i>I</i>) and Gasoline (<i>j</i>)	Heating Oil (<i>i</i>) and Gasoline (<i>j</i>)
Contracts	F1, F4, F7, F11, F15 and F18	F1, F3, F5, F7 and F9	F1, F3, F5, F7 and F9
Period	09/09/1996 to 05/25/2009	06/30/1997 to 05/25/2009	06/30/1997 to 05/25/2009
Number obs.	664	622	622
μ_{ξ_i}	0.1048*** (0.0378)	0.1159*** (0.0427)	0.1184*** (0.0297)
μ_{ξ_j}	0.1248* (0.0372)	0.1479*** (0.0468)	0.1330*** (0.0173)
κ_i	1.1549*** (0.0163)	1.5571*** (0.0323)	2.0123*** (0.0547)
κ_j	1.5461*** (0.0147)	2.0262*** (0.0758)	2.2203*** (0.0573)
σ_{ξ_i}	0.1824*** (0.0038)	0.2434*** (0.0056)	0.2077*** (0.0044)
σ_{ξ_j}	0.1741*** (0.0037)	0.2740*** (0.0094)	0.2374*** (0.0056)
σ_{χ_i}	0.3091*** (0.0075)	0.2577*** (0.0049)	0.2672*** (0.0064)
σ_{χ_j}	0.3051*** (0.0086)	0.4047*** (0.0117)	0.3809*** (0.0088)
λ_{ξ_i}	0.0901** (0.0377)	0.2931*** (0.0427)	-0.0194 (0.0305)
λ_{ξ_j}	0.0827** (0.0370)	0.1551*** (0.0505)	0.2115*** (0.0179)
λ_{χ_i}	0.1716** (0.0633)	0.2491*** (0.0524)	-0.0534 (0.0533)
λ_{χ_j}	0.4984** (0.0697)	0.4998*** (0.0768)	0.4999*** (0.0001)
φ_i			1.0035*** (0.0003)
φ_j	0.9955*** (0.0002)	1.0033*** (0.0003)	0.9994*** (0.0005)
$\rho_{\xi_i \xi_j}$	0.8774*** (0.0091)	0.9424*** (0.0049)	0.6803*** (0.0220)
$\rho_{\xi_i \chi_i}$	-0.0679** (0.0335)	-0.3623*** (0.0298)	0.1512*** (0.0362)
$\rho_{\xi_i \chi_j}$	0.0859*** (0.0349)	-0.1849*** (0.0437)	0.2752*** (0.0279)
$\rho_{\xi_j \chi_i}$	0.2304*** (0.0314)	-0.2486*** (0.0507)	0.5084*** (0.0213)
$\rho_{\xi_j \chi_j}$	0.1356*** (0.0355)	-0.1036 (0.0675)	-0.2050*** (0.0338)
$\rho_{\chi_i \chi_j}$	0.7125*** (0.0192)	0.8328*** (0.0073)	0.0458 (0.0345)
σ_η	0.0142*** (0.0001)	0.0125*** (0.0001)	0.0142*** (0.0001)
Log-likelihood	54209.19	42004.52	40908.09
AIC	54169.19	41964.52	40868.09
SIC	54079.22	41875.86	40779.43

The table presents the results for the Schwartz and Smith (2000) two-factor model without a common long-term trend for pairs of commodities. Standard errors are in parentheses. The estimated values are reported with * denoting significance at 10%, ** denoting significance at 5%, and *** denoting significance at 1%.

Second, in the specifications with a common long-term trend (Tables 6.5 and 6.7), the values of the constants C_i are very close to one, suggesting that the common long-term trend affects all three commodities in the same way and that the price-level differences between crude oil and refined products are stationary. This fact confirms the previous finding that the refining margin is stationary and, therefore, exhibits different dynamics than crude oil or refined products. In addition, it suggests that crude oil, heating oil and gasoline prices are not only cointegrated but also share a common long-term trend.

Moreover, the values of the K_i constants in Tables 6.5 and 6.7 are positive, confirming the differences in price levels among crude oil and refined products observed in Table 6.1. Of the three commodities considered, the most expensive is gasoline, followed by heating oil and crude oil.

If we define the Schwartz Information Criterion (SIC) as $\ln(L_{ML}) - q \ln(T)$, where q is the number of estimated parameters, T is the number of observations and L_{ML} is the value of the likelihood function using the q estimated parameters, then the fit is better when the SIC is higher. The same conclusions are obtained with the Akaike Information Criterion (AIC), which is defined as $\ln(L_{ML}) - 2q$.

It is worth noting that in Table 6.3, the values of the SIC and the AIC are higher for crude oil than for the other two commodities within the same time period (06/30/1997 to 05/25/2009) and the same

TABLE 6.7 The Joint Model With a Common Long-Term Trend for Pairs of Commodities

	WTI Crude Oil (<i>i</i>) and Heating Oil (<i>j</i>)	WTI Crude Oil (<i>i</i>) and Gasoline (<i>j</i>)	Heating Oil (<i>i</i>) and Gasoline (<i>j</i>)
Contracts	F1, F4, F7, F11, F15 and F18	F1, F3, F5, F7 and F9	F1, F3, F5, F7 and F9
Period	09/09/1996 to 05/25/2009	06/30/1997 to 05/25/2009	06/30/1997 to 05/25/2009
Number obs.	664	622	622
μ_{ξ}	0.1053** (0.0425)	0.1238** (0.0534)	0.1321 *** (0.0331)
κ_i	1.1689 *** (0.0142)	1.1804*** (0.0338)	1.3344*** (0.0433)
κ_j	1.3696*** (0.0164)	1.1047*** (0.0287)	1.0773*** (0.0281)
σ_{ξ}	0.2019*** (0.0187)	0.2393*** (0.0069)	0.2114*** (0.0052)
σ_{χ^i}	0.7989*** (0.0126)	0.3377*** (0.0093)	0.3443*** (0.0105)
σ_{χ^j}	0.6080*** (0.0001)	0.3462*** (0.0092)	0.3683*** (0.0110)
λ_{ξ}	0.0718** (0.0416)	0.0571 * (0.0328)	0.0439* (0.0249)
λ_{χ^i}	0.4589*** (0.0707)	-0.1255 (0.0997)	-0.0111 (0.0617)
λ_{χ^j}	0.2463*** (0.0810)	-0.7679*** (0.0734)	-0.4750*** (0.0495)
φ_i			1.0016*** (0.0003)
φ_j	0.9944*** (0.0002)	1.0036*** (0.0002)	1.0011 *** (0.0006)
$\rho_{\xi\chi^i}$	-0.3323* (0.1904)	-0.3753*** (0.0347)	-0.2651 *** (0.0414)
$\rho_{\xi\chi^j}$	-0.2322 (0.1861)	-0.2047*** (0.0415)	-0.2340*** (0.0456)
$\rho_{\chi^i\chi^j}$	0.9532*** (0.0001)	0.8137*** (0.0184)	0.8242*** (0.0167)
K_i			
K_j	0.5594*** (0.1312)	-1.6122*** (0.2819)	1.6669*** (0.2270)
C_i			
C_j	1.0025*** (0.0034)	0.8862*** (0.0061)	0.8932*** (0.0049)
σ_{η}	0.0161 *** (0.0001)	0.0165*** (0.0001)	0.0174*** (0.0001)
Log-likelihood	53061.99	41007.69	40517.87
AIC	53029.99	40975.69	40483.87
SIC	52958.02	40904.76	40408.51

The table presents the results for the Schwartz and Smith (2000) two-factor model assuming a common long-term trend for pairs of commodities. Standard errors are in parentheses. The estimated values are reported with * denoting significance at 10%, ** denoting significance at 5%, and *** denoting significance at 1%.

contracts (F1, F3, F5, F7 and F9).^{*} The superior fit for the crude oil data can be related to the short-term imperfections in the refined product market and, particularly in the case of gasoline, the lack of liquidity in longer-maturity futures price quotes. This result is to be expected because crude oil has futures with much more liquidity at longer maturities than refined products. In the case of refined products, particularly gasoline, maturities above several months are not as liquid as would be required to calibrate such models.

The relative fit of the models to our three commodity price series can be assessed by evaluating their predictive ability. The in-sample predictive ability of the Schwartz and Smith (2000) two-factor model is presented in Table 6.8 for the three specifications considered in Section 6.3.1. We found that the differences among the three models, in terms of bias and the root mean squared error, are small. Moreover, in Table 6.8, it can be observed that the root mean squared error values for crude oil are generally lower than the corresponding values obtained for the other two commodities, confirming our previous prediction.

The out-of-sample predictive ability of the Schwartz and Smith (2000) two-factor model for the three specifications is presented in Table 6.9. The results were obtained by valuing the contracts corresponding

* It should be pointed out that although in principle the values of the SIC and AIC are not directly comparable whenever the series are different, our three commodity price series have the same order of magnitude; the time period (06/30/1997 to 04/24/2006) and the futures contracts used in the estimation procedure (F1, F3, F5, F7 and F9) are the same.

TABLE 6.8 In-Sample Predictive Ability.

Panel A. Model for the Commodities Separately

WTI			Heating Oil			Gasoline		
Contract	ME	RMSE	Contract	ME	RMSE	Contract	ME	RMSE
F1	−0.0005	0.0455	F1	0.0154	0.047	F1	−0.0053	0.0509
F3	0.0013	0.0394	F3	0.0159	0.0438	F3	−0.004	0.0436
F5	−0.0004	0.0353	F5	0.0144	0.0398	F5	−0.0039	0.0381
F7	−0.0009	0.0323	F7	0.0126	0.0344	F7	−0.0038	0.0345
F9	0.0006	0.0309	F9	0.0119	0.0327	F9	−0.0023	0.0337

Panel B. Joint Model Without a Common Long-Term Trend

WTI			Heating oil			Gasoline		
Contract	ME	RMSE	Contract	Mean	RE	Contract	ME	RMSE
F1	−0.0053	0.0463	F1	0.0021	0.0452	F1	−0.0068	0.0476
F3	−0.0028	0.0386	F3	0.0016	0.0389	F3	−0.005	0.0396
F5	−0.0031	0.0325	F5	0.0013	0.0341	F5	−0.0033	0.0329
F7	−0.0031	0.0284	F7	−0.0006	0.028	F7	−0.0033	0.0285
F9	−0.0016	0.0266	F9	−0.001	0.0263	F9	−0.0028	0.0274

Panel C. Joint Model With a Common Long-Term Trend

WTI			Heating Oil			Gasoline		
Contract	ME	RMSE	Contract	ME	RMSE	Contract	ME	RMSE
F1	−0.0103	0.049	F1	−0.0121	0.0551	F1	−0.0033	0.0467
F3	−0.0069	0.039	F3	−0.0089	0.0442	F3	−0.0033	0.0381
F5	−0.0082	0.033	F5	−0.0068	0.0391	F5	−0.0031	0.0336
F7	−0.0105	0.03	F7	−0.0077	0.0337	F7	−0.0043	0.0315
F9	−0.0125	0.0292	F9	−0.0083	0.0309	F9	−0.005	0.0321

The table presents the mean error (ME), calculated as real minus predicted values, and the root mean squared error (RMSE), to analyse the in-sample predictive power ability of the Schwartz and Smith (2000) two-factor model for the three commodities separately (panel A), and for all three commodities both without a common long-term trend (panel B) and with a common long-term trend (panel C). The time period is 06/30/1997 to 05/25/2009 (622 weekly observations for each commodity).

to the period 06/16/2003 to 05/25/2009 with the parameters obtained estimating the models with data from 06/30/1997 to 06/09/2003 (311 observations in each period). As expected, out-of-sample pricing errors are slightly higher than the corresponding in-sample values. As before, the differences between the three models are small.

As a result, given that the predictive ability of the two models is very similar, we can conclude that all three commodities exhibit the same long-term trend. Consequently, because the joint model that assumes common long-term trends is the simplest one, we do not need a second long-term factor when jointly modelling the three commodities.

It is also worth noting that although previous studies have already found that the prices of crude oil and its main refined products are cointegrated, this conclusion is extended in the present work because we find that they also exhibit a common long-term trend. Moreover, to the best of our knowledge, this is the first time that a factor model with a common long-term trend for the prices of crude oil and its main refined product has been proposed and estimated. Additionally, the results of the estimation of this common long-term trend model suggest that the price-level differences between crude oil and refined products are stationary, confirming that the refining margin is stationary and therefore, exhibits different dynamics from those of crude oil and refined products.

TABLE 6.9 Out-of-Sample Predictive Ability

Panel A. Model for the commodities separately

WTI			Heating Oil			Gasoline		
Contract	ME	RMSE	Contract	ME	RMSE	Contract	ME	RMSE
F1	−0.003	0.0456	F1	−0.0062	0.045	F1	−0.0045	0.0556
F3	0.0015	0.0402	F3	0.0036	0.0438	F3	−0.0044	0.0472
F5	0.0004	0.0379	F5	0.0043	0.0401	F5	−0.0044	0.0429
F7	0.0001	0.0363	F7	0.0021	0.0354	F7	−0.003	0.0422
F9	0.0016	0.0354	F9	−0.0018	0.0414	F9	0.0001	0.0535

Panel B. Joint Model Without a Common Long-Term Trend

WTI			Heating Oil			Gasoline		
Contract	ME	RMSE	Contract	Mean	RE	Contract	ME	RMSE
F1	−0.0085	0.0471	F1	0.0002	0.0448	F1	−0.0073	0.058
F3	−0.0029	0.0404	F3	0.004	0.0435	F3	−0.0049	0.0477
F5	−0.003	0.0379	F5	0.0024	0.0396	F5	−0.0044	0.0438
F7	−0.0026	0.036	F7	0.001	0.0355	F7	−0.003	0.0423
F9	−0.0004	0.0347	F9	−0.0001	0.0448	F9	−0.0002	0.0571

Panel C. Joint Model With a Common Long-Term Trend

WTI			Heating Oil			Gasoline		
Contract	ME	RMSE	Contract	ME	RMSE	Contract	ME	RMSE
F1	−0.0192	0.0528	F1	−0.0365	0.0622	F1	0.021	0.0572
F3	−0.0073	0.0411	F3	−0.0103	0.049	F3	0.0016	0.0477
F5	−0.0049	0.0388	F5	0.0054	0.0424	F5	−0.0166	0.0516
F7	−0.0047	0.0383	F7	0.0171	0.0399	F7	−0.0313	0.058
F9	−0.005	0.0388	F9	0.0254	0.0471	F9	−0.0422	0.0657

The table presents the mean error (ME), calculated as real minus predicted values, and the root mean squared error (RMSE) to analyse the out-of-sample predictive power ability of the Schwartz and Smith (2000) two-factor model for the three commodities separately (panel A) and for all three commodities jointly without a common long-term trend (panel B) and with a common long-term trend (panel C). The results are obtained valuing the contracts corresponding to the period 06/16/2003 to 05/25/2009 using the parameters obtained by estimating the models with data from 06/30/1997 to 06/09/2003 (311 observations in each period).

Finally, the fact that crude oil and refined product prices share a common long-term trend has important implications for managing and hedging the refining margin risk. With a single long-term trend, the refining margin reflects only short-term effects. However, if we assume different long-term trends for each commodity, the refining margin reflects long- and short-term effects, implying higher volatility and therefore, a higher value of derivatives for the refining margin.

In the next section, we will use these results to value the main products used to protect the refining margin, i.e. the crack-spread options listed by NYMEX.

6.4 Crack-Spread Option Valuation

It is well known that commodity markets have grown quickly during recent years. Four types of investors can be distinguished in commodity markets: investors seeking a portfolio diversification tool, particularly pension funds; investors seeking a source of alpha, particularly hedge funds; European banks, which use commodity derivatives to structure products that they then retail to their customers; and energy companies protecting their profits.

Energy companies include refining companies that are willing to hedge their refining margin with specific hedging products. Crack-spread options are the most popular products for this purpose. As stated above, a crack-spread call option is a contract that gives the holder the right (but not the obligation) to buy a refined product futures contract from the writer and sell the writer a crude oil futures contract, paying a previously agreed-upon crack-spread price. A crack-spread put allows one to sell a refined product futures contract and buy a crude oil futures contract, paying a crack-spread price.

In valuing this type of product, it is essential to consider our previous results. In particular, it is crucial to realize that the refining margin dynamics are completely different from those of crude oil or refined products. As shown above, the price dynamics of crude oil and refined products are integrated, whereas the refining margin dynamics are stationary. Even more importantly, crude oil and refined products are not only cointegrated but also exhibit a common long-term trend; as we will observe, this finding is of critical importance in managing and hedging refining margin risk, primarily in the valuation of crack-spread options. Specifically, the values of crack-spread options must be lower if we assume a common long-term trend for crude oil and refined products than if the refining margin is integrated because in the case of a common long-term trend, the refining margin risk reflects only short-term effects. Therefore, there is less volatility and the writer of the crack-spread option faces lower risk than in cases where there are different long-term trends for each commodity. Consequently, we conclude that we must consider a refining margin model that only includes short-term effects (i.e. a model in which crude oil and its main refining products share the long-term factors).

In this section, we compare the pricing performance of the model with a common long-term trend, the model that allows for a long-term trend for each commodity, the model that postulates uncorrelated sub-models for each commodity, and the approach by Dempster *et al.* (2008) for directly modelling the crack-spread.* We show that the most suitable method of valuing crack-spread options is to assume a common long-term trend for WTI crude oil and gasoline and for WTI and heating oil.

6.4.1 Data

The data set used in the estimation procedure consists of two sets of daily observations of crack-spread put and call options quoted at the NYMEX. The first set consists of heating oil vs. WTI data, and the second one consists of gasoline vs. WTI data.

In the NYMEX, there are only quotations for these two types of crack-spread options (i.e. heating oil vs. WTI and gasoline vs. WTI), and contracts mature each month for the following 18 months at different strike prices. Specifically, the strike prices are the one at-the-money strike price, five additional strikes both above and below the established at-the-money strike price in \$0.25 (25¢) increments, three additional out-of-the-money strike prices above and below those strikes at \$1.00 intervals and two more strikes added above and below at \$2.00 intervals. Options traded at NYMEX are American-style, and thus, the holder can exercise his or her right at any time.

The market for this type of option is much less liquid than the market for futures. Due to the scarcity of data, we have chosen to use daily rather than weekly data. As described below, this has forced us to make some minor changes to the model. Specifically, for the case of gasoline vs. WTI crack-spread options, we have daily data from November 2001 to May 2009 (20,243 observations), with exercise prices ranging from \$3 to \$17. In the case of the heating oil vs. WTI crack-spread options, we have daily data corresponding to contracts maturing from August 2000 to November 2009 (6663 observations), with exercise prices from \$3 to \$14. A brief summary of these options is given in [Table 6.10](#).

* A short summary of the Dempster *et al.* (2008) model is presented in Appendix B.

TABLE 6.10 Descriptive Statistics for Crack-Spread Options

	% Sample	Mean Number of Observations Series	Mean K (\$)	Mean Price (\$)
Panel A: Heating oil vs. WTI (169 series, 6663 observations)				
Put options	74.66	42.86	8.24	1.0111
Call options	25.44	29.37	8.67	1.3044
Panel B: Gasoline vs. WTI (188 series, 20,243 observations)				
Put options	67.55	143.66	8.50	1.1647
Call options	32.45	32.75	6.03	2.0986

The table shows the main descriptive stats of the crack-spread options. The time period is November 2001 to May 2009 for gasoline vs. WTI crack-spread options, and there are exercise prices from 3 to 17 dollars. For heating oil vs. WTI crack-spread options, there are data corresponding to contracts maturing from August 2000 to November 2009, with exercise prices from 3 to 14 dollars.

6.4.2 Option Valuation Methodology

In this study, we have a set of crack-spread American options that we wish to replicate. There is no close analytic expression of their price, so a simulation must be used. In addition, our models are either three-dimensional (i.e. the common trended model) or four-dimensional (i.e. when we consider one trend for each commodity), severely narrowing the methods we can use.

There are several methods for valuing American options. Each method involves the discretization of the state space, such that at each point in time and for every value of the factors, the option holder must decide whether to exercise his right; therefore, the decision can only be based on future dynamics. However, when the number of factors increases, one encounters the ‘curse of dimensionality’ (a concept conceived by Bellman) because full discretization is almost infeasible due to the computational complexity inherent in there being more than three factors.

To minimize this problem, one of three main approaches can be selected (Bally *et al.* 2005). One approach is to perform state aggregation to develop a ‘synthetic indicator’ such that the holder of the option can decide whether to exercise it using only this indicator rather than three or more factors (Barraquand and Martineau 1995). The second approach computes conditional expectations via Malliavin Calculus.* The third one uses base functions (Tsitsiklis and VanRoy 1999, Longstaff and Schwartz 2001). Among all these approaches, the one by Longstaff and Schwartz (2001), called Least Squares Monte Carlo, has become one of the most popular models in the literature for valuing American options and will be the approach used in this paper. However, the point here is not the specific method used to value the crack-spread options but rather whether the valuation results are better when a common long-term trend is assumed than when different long-term trends are assumed.

The parameters and dynamics needed for option valuation, for the joint models with and without a common long-term trend and for the uncorrelated model, were estimated in Section 6.3 using weekly data. Specifically, given that crack-spread options involve two commodities, we used the results in Tables 6.6 and 6.7.† However, changing the model to daily dynamics is not a problem in continuous time. We simply re-sampled the model using an aliasing algorithm to obtain information on daily states. The RMSE is slightly higher, though the difference is small.

* See Fournié *et al.* (1993) for numerical applications and Bouchard and Touzi (2004) for general theory and variance reduction.

† As stated in Appendix B, the model by Dempster *et al.* (2008) is estimated using the same procedure (the Kaman filter) and the same data sets as the rest of the models presented in Section 6.3.

6.4.3 Commodity Price Dynamics

For the purpose of option valuation, we need a full description of the model. In matrix form, the state dynamics can be described as follows:

$$dZ_t = (\mu + AZ_t) dt + dW_t. \quad (6.5)$$

To clarify matters, let U_t be a unit of Brownian motion (i.e. $dU_t dU_t^T = I dt$) and rewrite (6.5) as

$$dZ_t = (\mu + AZ_t) dt + R dU_t. \quad (6.6)$$

The models used for crack-spread option valuation will be the same as those proposed in [Section 6.3](#) for the commodities in pairs: the model for two commodities with no correlation (i.e. the first specification), the joint model without a common long-term trend (i.e. the second specification), the joint model with a common long-term trend (i.e. the third specification), together with the model by Dempster *et al.* (2008).

In the case where two models are combined, we have two equations in matrix form:

$$dZ_{it} = (\mu_i + A_i Z_{it}) dt + R_i dW_{it} \quad (i = 1, 2). \quad (6.7)$$

Therefore, if we assume no correlation, the global model is

$$\begin{pmatrix} dZ_{1t} \\ dZ_{2t} \end{pmatrix} = \left[\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix} \begin{pmatrix} Z_{1t} \\ Z_{2t} \end{pmatrix} \right] dt + \begin{pmatrix} R_1 & 0 \\ 0 & R_2 \end{pmatrix} \begin{pmatrix} dU_{1t} \\ dU_{2t} \end{pmatrix}, \quad (6.8)$$

which is equivalent to estimating both models separately. If we allow a free correlation structure, the model is

$$\begin{pmatrix} dZ_{1t} \\ dZ_{2t} \end{pmatrix} = \left[\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix} \begin{pmatrix} Z_{1t} \\ Z_{2t} \end{pmatrix} \right] dt + R^* \begin{pmatrix} dU_{1t} \\ dU_{2t} \end{pmatrix}, \quad (6.9)$$

where R^* is a free triangular matrix.

For parameter estimation purposes (see [Appendix A](#)), we use Kalman filter equations to estimate $Z_t|_{t-1} = E[Z_t/Z_1, \dots, Z_{t-1}]$, and as an intermediate result, $Z_{t-1}|_{t-1} = E[Z_{t-1}/Z_1, \dots, Z_{t-1}]$. This process (estimating using current or even future information) is termed ‘aliasing’ in the Kalman filter literature, and the series $Z_{t|t}$ will be used as initial states for option valuation.

6.4.4 Results

[Table 6.11](#) presents several metrics to analyse the in-sample predictive ability of the models for heating oil vs. WTI and gasoline vs. WTI crack-spread options. The models considered are the (three-factor) joint model with a common long-term trend, the (four-factor) joint model without a common long-term trend, the (two-factor) model for both commodities separately and the model by Dempster *et al.* (2008).

The statistics presented in [Table 6.11](#) are the mean bias (real minus predicted), the root mean squared error, the bias standard deviation and the mean absolute error. The values shown in the table are the mean of the different means or standard deviations for each option series.

It can be observed that we achieve the best results with the common long-term trend model (i.e. the three-factor model). The common long-term trend model achieves even better results than the one by

TABLE 6.11 Crack-Spread Option Valuation Results. Error Descriptive Statistics

Statistic	Three-Factor Model (Joint Model with Common Trend)	Four-Factor Model (Joint Model Without Common Trend)	Two-Factor Model (Uncorrelated Model)	Two-Factor Model For The Spread
Panel A: Heating oil vs. WTI crack-spread options				
Mean bias (real – predicted)	–1.4268	–3.1037	–3.0287	–2.6674
RMSE	1.1859	3.4725	3.2371	3.1717
Bias standard deviation	0.3982	0.4376	0.4381	0.3352
Mean absolute error	1.7378	3.3738	3.1429	3.0988
Panel B: Gasoline vs. WTI crack-spread options				
Mean bias (real – predicted)	–0.7465	–1.3125	–1.3605	–1.6759
RMSE	1.3637	1.6608	1.6884	2.3714
Bias standard deviation	0.5190	0.5566	0.5847	1.0432
Mean absolute error	1.2173	1.4912	1.5148	1.9204

The table presents several metrics to analyse the in-sample predictive power ability of the models under study: the joint model with a common trend, the joint model without a common trend, the model tracking commodities separately and the two-factor model for the crack-spread by Dempster *et al.* (2008). The time period is November 2001 to May 2009 for gasoline vs. WTI crack-spread options (20,243 daily observations) and August 2000 to November 2009 for heating oil vs. WTI crack-spread options (6663 daily observations). For each series, we have used the corresponding statistic. These results correspond to the mean of these multiple means or standard deviations.

Dempster *et al.* (2008). The uncorrelated model achieves the worst results, thus, confirming our hypothesis regarding common trending. It is worth noting that, as expected, the theoretical option values are lower (and closer to the real values) when we assume a common long-term trend than when we allow for different long-term trends. This result may be related to the fact that there is less volatility and that the writer of the option assumes lower risk when there is a common long-term trend for all three commodities. The fact that all three commodities share a common long-term trend implies that the refining margin risk reflects only short-term effects; therefore, there is less volatility, and the writer of the option faces lower risk than in the case when there are different long-term trends for each commodity. With different long-term trends for each commodity, the refining margin reflects long- and short-term effects, implying higher volatility; therefore, crack-spread option values are higher.

The better performance of the common long-term trend model compared to the model by Dempster *et al.* (2008) also has important implications. As explained in the [Section 6.1](#), the approach by Dempster *et al.* (2008) is advantageous because by directly modelling the spread, we can avoid modelling the correlation between two asset returns, which is difficult (Kirk 1995, Mbanefo 1997, Alexander 1999). On the other hand, the common long-term trend model proposed in this paper is advantageous because, using two or more price series, it is possible to extract more information for estimation purposes (Cortazar *et al.* 2008). *A priori*, it is difficult to know which of the two approaches gives better valuation results. In this paper, we show that at least with our database, the common long-term trend model gives better valuation results in terms of option pricing errors.

Finally, we can conclude that because the valuation errors obtained using the joint model with a common long-term trend are smaller than those obtained using the joint model without a common long-term trend or the uncorrelated model, the preferred model is the one that assumes a common long-term trend, which confirms the [Section 6.3](#) findings. This result also implies that the advantage of the model that has a long-term trend in common (i.e. the possibility of extracting more information for estimation purposes using two or more price series) is more important with respect to valuation errors than the advantage of the Dempster *et al.* (2008) model (i.e. avoiding the need to model the correlation between two asset returns).

This result also provides additional evidence suggesting that the refining margin exhibits dynamics that are different from those of crude oil and refined products, which has crucial implications not only in terms of crack-spread option valuation but also in terms of managing and hedging the refining margin risk.

6.5 Conclusion

In this paper, we show that, in contrast to crude oil and its main refined products (i.e. gasoline and heating oil), the refining margin is stationary. Moreover, we found evidence suggesting that these three commodities are not only cointegrated, as shown in previous papers, but also share common long-term dynamics, suggesting that the refining margin only reflects short-term effects.

We confirmed the hypothesis of a common long-term trend for crude oil, heating oil and gasoline by proposing different factor models to jointly estimate the dynamics of these three commodities. These factor models were framed within the multi-factor model family proposed by Schwartz (1997) and a related series of papers by Schwartz and Smith (2000), Cortazar and Schwartz (2003), Cortazar and Naranjo (2006) and Cortazar *et al.* (2008), among others.

Furthermore, we found that among different factor models with and without common long-term trends, the most suitable model in terms of simplicity and fit is the one that assumes a common long-term trend for all three commodities.

The fact the these three commodities share a common long-term trend has important implications in terms of managing and hedging the risk faced by a refining companies; specifically, it has important implications in the valuation of crack-spread options, which are used to protect the refining margin. Given that these crack-spread options involve several assets, the techniques used for their valuation are complex. In this paper, we present a simplified valuation model for this type of option, which assumes a common long-term trend for crude oil, heating oil and gasoline.

We found that by assuming a common long-term trend for the prices of crude oil and refined products, option valuation was more accurate than using models with more factors and parameters, including the model by Dempster *et al.* (2008), which directly models the spread. Therefore, we conclude that the preferred model for valuing this kind of option is one that assumes a common long-term trend; such a model not only yields better results in terms of pricing errors but is also simpler and easier to implement than standard techniques. Moreover, these results point out that using two or more price series, it is possible to extract more information for estimation purposes, which has more importance in terms of valuation errors than avoiding the need to model the correlation between two asset returns.

In summary, the primary finding of this paper is that the refining margin exhibits dynamics that are different from those of crude oil and its main refined products because the refining margin risk only reflects short-term effects. This finding has important implications for crack-spread option valuation, and more generally, it also has many other implications for the process of managing and hedging the risks faced by refining companies.

References

- Alexander, C., Correlation and cointegration in energy markets. In *Managing Energy Price Risk*, 2nd ed., edited by V. Kaminsky, pp. 573–606, 1999 (Risk Publications: London).
- Asche, F., Gjolberg, O. and Völker, T., Price relationships in the petroleum market: an analysis of crude oil and refined product prices. *Energy Econ.*, 2003, **25**, 289–301.
- Bally, V., Caramellino, L. and Zanette, A., Pricing American options by Monte Carlo methods using a Malliavin Calculus approach. *Monte Carlo Meth. Applic.*, 2005, **11**, 97–133.
- Barraquand, J. and Martineau, D., Numerical valuation of high dimensional multivariate American securities. *J. Financ. Quantit. Anal.*, 1995, **30**(3), 383–405.
- Bouchard, B. and Touzi, N., On the Malliavin approach to Monte Carlo approximation of conditional expectations. *Finance Stochast.*, 2004, **8**(1), 45–71.
- Cortazar, G., Milla, C. and Severino, F., A multicommodity model for futures prices: using futures prices of one commodity to estimate the stochastic process of another. *J. Fut. Mkts.*, 2008, **28**(6), 537–560.
- Cortazar, G. and Naranjo, L., An N-factor Gaussian model of oil futures prices. *J. Fut. Mkts.*, 2006, **26**(3), 209–313.

- Cortazar, G. and Schwartz, E.A., Implementing a stochastic model for oil futures prices. *Energy Econ.*, 2003, **25**, 215–218.
- Dempster, M.E., Medova, E. and Tang, K., Long term spread option valuation and hedging. *J. Bank. Finance*, 2008, **32**, 2530–2540.
- Fournié, E., Lasry, J.M., Lebouchoux, J., Lions, P.L. and Touzi, N., Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance Stochast.*, 1993, **3**, 391–412.
- Froot, K.A., Schwartstein, D.S. and Stein, J.C., Risk management: coordinating corporate investment and financial policies. *J. Finance*, 1993, **48**, 1629–1658.
- García, A., Población, J. and Serna, G., The stochastic seasonal behavior of natural gas prices. *Eur. Financ. Mgmt.*, 2012, **18**, 410–443.
- Gjolberg, O. and Johnsen, T., Risk management in the oil industry: can information on long-run equilibrium prices be utilized? *Energy Econ.*, 1999, **21**(6), 517–527.
- Graham, J.R. and Smith, C.W., Tax incentives to hedge. *J. Finance*, 1999, **54**, 2241–2262.
- Holmström, B. and Tirole, J., Liquidity and risk management. *J. Money, Credit Bank*, 2000, **32**, 295–319.
- International Energy Agency. Oil Market Report. Annual Statistical Supplement, 2006. Available online at: <http://www.omrpublic.iea.org>
- Kirk, E., Correlation in the energy markets. In: *Management Energy Price Risk*, pp. 177–186, 1995 (Risk Publications: London).
- Lanza, A., Manera, M. and Giovannini, M., Modeling and forecasting cointegrated relationships among heavy oil and product prices. *Energy Econ.*, 2005, **27**, 831–848.
- Longstaff, F. and Schwartz, E., Valuing American options by simulations: a simple least squares approach. *Rev. Financ. Stud.*, 2001, **14**(1), 113–147.
- Mayers, D., Smith, J. and Clifford, W., On the corporate demand for insurance. *J. Business*, 1982, **55**, 281–296.
- Mbanefo, A., Co-movement term structure and the valuation of energy spread options. In *Mathematics of Derivative Securities*, edited by M.A.H. Dempster and R.R. Pliska, pp. 99–102, 1997 (Cambridge University Press: Cambridge).
- Mello, A.S. and Parsons, J.E., Hedging and liquidity. *Rev. Financ. Stud.*, 2000, **13**, 127–153.
- Pindyck, R.S., The long-run evolution of energy prices. *Energy J.*, 1999, **20**(2), 1–27.
- Schwartz, E.S., The stochastic behavior of commodity prices: implication for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Schwartz, E.S. and Smith, J.E., Short-term variations and long-term dynamics in commodity prices. *Mgmt Sci.*, 2000, **46**(7), 893–911.
- Serletis, A., Unit root behavior in energy futures prices. *Energy J.*, 1992, **13**, 119–128.
- Serletis, A., A cointegration analysis of petroleum futures prices. *Energy Econ.*, 1994, **16**(2), 93–97.
- Smith, C.W. and Slutz, R.M., The determinants of firms hedging policies. *J. Financ. Quantit. Anal.*, 1985, **20**, 391–405.
- Sorensen, C., Modeling seasonality in agricultural commodity futures. *J. Fut. Mkts*, 2002, **22**, 393–426.
- Tsitsiklis, J.N. and VanRoy, B., Optimal stopping for Markov processes: Hilbert space theory, approximation algorithm and an application to pricing high-dimensional financial derivatives. *IEEE Trans. Autom. Control*, 1999, **44**, 1840–1851.

Appendix A: Estimation Methodology

The Kalman filter technique is a recursive methodology that estimates the unobservable time series, the state variables or the factors (Z_t) based on an observable time series (Y_t) that depends on these state variables. The *measurement equation* accounts for the relationship between the observable time series and the state variables:

$$Y_t = dt + MtZ_t + \eta_t, t = 1, \dots, Nt, \quad (\text{A.1})$$

where $Y_t, d_t \in \mathbb{R}^n$, $M_t \in \mathbb{R}^{n \times m}$, $Z_t \in \mathbb{R}^h$, h is the number of state variables, or factors, in the model, and $\eta_t \in \mathbb{R}^n$ is a vector of serially uncorrelated Gaussian disturbances with zero mean and covariance matrix H_t .

In the estimation procedure, a discrete time version of this equation is necessary; in the case of the joint model with a common long-term trend for the three commodities, this equation is given by the following expressions:

$$Y_t = \begin{pmatrix} \ln F_{T_1}^1 \\ \vdots \\ \ln F_{T_n}^1 \\ \ln F_{T_1}^2 \\ \vdots \\ \ln F_{T_n}^2 \\ \ln F_{T_1}^3 \\ \vdots \\ \ln F_{T_n}^3 \end{pmatrix}, dt = \begin{pmatrix} A^1(T_1) \\ \vdots \\ A^1(T_n) \\ A^2(T_1) \\ \vdots \\ A^2(T_n) \\ A^3(T_1) \\ \vdots \\ A^3(T_n) \end{pmatrix}$$

$$M_t = \begin{pmatrix} 1 & e^{-k_1 T_1} & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & e^{-k_1 T_n} & 0 & 0 \\ 1 & 0 & e^{-k_2 T_1} & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & e^{-k_2 T_n} & 0 \\ 1 & 0 & 0 & e^{-k_3 T_1} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & e^{-k_3 T_n} \end{pmatrix},$$

and $F_{T_1}^i$ is the price of a futures contract for the commodity 'i' ($i = 1, 2, 3$) with maturity at time ' $T_1 + t$ ' traded at time t . In principle, it would be possible to use a different number of futures contracts for each commodity; however, in this work, we consider it more suitable to use the same number (' n ') of futures contracts for all commodities.

The *transition equation* accounts for the evolution of the state variables:

$$Z_t = c_t + T_t Z_{t-1} + \psi_t, \quad t = 1, \dots, N_t, \quad (\text{A.2})$$

where $c_t \in \mathbb{R}^h$, $T_t \in \mathbb{R}^{h \times h}$ and $\psi_t \in \mathbb{R}^h$ is a vector of serially uncorrelated Gaussian disturbances with zero mean and covariance matrix Q_t .

In the case of the joint model with a common long-term trend for the three commodities, the discrete time version of this equation, which is needed in the estimation procedure, is given by the following expressions:

$$Z_t = \begin{pmatrix} \xi_{1t} \\ \chi_{1t} \\ \chi_{2t} \\ \chi_{3t} \end{pmatrix}, c_t = \begin{pmatrix} \mu_{\xi_1} \Delta t \\ 0 \\ 0 \\ 0 \end{pmatrix},$$

$$T_t = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-k_1 \Delta t} & 0 & 0 \\ 0 & 0 & e^{-k_2 \Delta t} & 0 \\ 0 & 0 & 0 & e^{-k_3 \Delta t} \end{pmatrix},$$

and

$$Var(\psi_t) = \begin{pmatrix} \sigma_{\xi_t}^2 \Delta t & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-k_1 \Delta t}) / k_1 & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} \Delta t & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-k_1 \Delta t}) / k_2 \\ \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-k_1 \Delta t}) / k_1 & \sigma_{\xi_t}^2 (1 - e^{-2k_1 \Delta t}) / (2k_1) & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_1 \Delta t + k_2 \Delta t)}) / (k_1 + k_2) & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_1 \Delta t + k_3 \Delta t)}) / (k_1 + k_3) \\ \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-k_1 \Delta t}) / k_2 & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_1 \Delta t + k_2 \Delta t)}) / (k_1 + k_2) & \sigma_{\xi_t}^2 (1 - e^{-2k_2 \Delta t}) / (2k_2) & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_2 \Delta t + k_3 \Delta t)}) / (k_2 + k_3) \\ \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-k_1 \Delta t}) / k_3 & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_1 \Delta t + k_3 \Delta t)}) / (k_1 + k_3) & \sigma_{\xi_t} \sigma_{\xi_t} \rho_{\xi_t \xi_t} (1 - e^{-(k_2 \Delta t + k_3 \Delta t)}) / (k_2 + k_3) & \sigma_{\xi_t}^2 (1 - e^{-2k_3 \Delta t}) / (2k_3) \end{pmatrix}$$

Here, $Y_t|_{t-1}$ is the conditional expectation of Y_t , and Ξ_t is the covariance matrix of Y_t conditional on all the information available at time $t - 1$. After omitting unessential constants, the log-likelihood function can be expressed as

$$l = - \sum_t \ln |\Xi_t| - \sum_t (Y_t - Y_{t|t-1})' \Xi_t^{-1} (Y_t - Y_{t|t-1}). \quad (\text{A.3})$$

Appendix B: Modelling the Crack-Spread Directly

Dempster *et al.* (2008) propose a model for valuing spread options on two commodity prices that are cointegrated. Let x_t denote the (latent) mean-reverting spot spread and y_t the mean-reverting process, with a zero long-run mean, representing the deviation from the long-term equilibrium of spot spreads. The expressions for x_t and y_t are*

$$\begin{aligned} dx_t &= k(\theta + \phi(t) + y_t - x_t) dt + \sigma dW, \\ dy_t &= -\kappa_2 y_t dt + \sigma_2 dW_2. \end{aligned} \quad (\text{B.1})$$

In this model, x and y revert towards θ and zero, with mean reversion speeds κ and κ_2 , respectively. dW and dW_2 are two standard Brownian motions with correlation ρ . $\phi(t)$ accounts for the seasonality effects in the spread, as in Dempster *et al.* (2008) and Sorensen (2002).

The model is estimated according to the Kalman filter methodology using two crack-spread series: the difference between heating oil and WTI crude oil prices and the difference between unleaded gasoline (RBOB) and WTI crude oil prices. In the case of heating oil vs. WTI crude oil, the data set used in the estimation comprises contracts F1, F4, F7, F11, F15 and F18, from 09/09/1996 to 05/25/2009 (664 weekly observations each contract). In the case of RBOB gasoline vs. WTI crude oil, the data set comprises contracts F1, F3, F5, F7 and F9, from 06/30/1997 to 05/25/2009 (622 weekly observations).

* To be consistent with our previous notation, the model by Dempster *et al.* (2008) is formulated here under the real measure.

Quantitative Spread Trading on Crude Oil and Refined Products Markets

7.1	Introduction	119
7.2	Crude Oil and Refined Products Markets	121
7.3	Optimal Statistical Arbitrage Trading Model.....	122
	Remark on Model Specification	
7.4	Multiple Hypothesis Testing: Data Snooping Bias	124
	Single-Step Procedure • Stepdown Procedure • Balanced Stepdown Procedure	
7.5	Outline of Empirical Analysis	129
	Remark on Transaction Costs • Remark on Liquidity	
7.6	Empirical Results	132
	Robustness to Transaction Costs	
7.7	Conclusion	135
	References.....	140
	Appendix A: Liquidity Measures.....	142
	Appendix B: Top 10 Profitable Trading Strategies	144

Mark Cummins

Andrea Bucca

Quantitative trading in oil-based markets is investigated over 2003–2010, with a focus on WTI, Brent, heating oil and gas oil. A total of 861 spreads are considered. A novel optimal statistical arbitrage trading model is applied, with generalised stepwise procedures controlling for data snooping bias. Aggregating upward and downward mean-reversion, profitable strategies are identified with Sharpe ratios greater than 2 in many instances. For the top categories, average daily returns range from 0.07 to 0.55%, with trade lengths of 9–55 days. A collapse in the number of profitable trading strategies is seen in 2008. Robustness to varying transactions costs is examined.

Keywords: Econometrics; Energy derivatives; Trading systems; Quantitative trading strategies

JEL Classification: C1, C5, C6, C12, C52, C63

7.1 Introduction

A significant increase in commodities trading has been observed over the past decade. This increased activity has emerged in a range of forms, including commodity index investment and the proliferation of commodity-based hedge funds. Mou (2010) reports that, by mid-2008, the value of long positions held by index investors reached \$256 bn, with the market ratio of index investment positions to total open interest at 44%. The collapse in commodity markets resulting from the credit crisis and global

economic recession concerns led to a huge sell off of positions, with the total index investment value falling to \$112 bn and the market ratio falling to 39%, although both measures had recovered by the end of 2009 to \$211 bn and 52%, respectively. Gregeriou *et al.* (2009) report that assets under management by the Commodity Trading Advisor (CTA) industry reached \$208 bn by the end of 2008. This increased commodity trading activity has led to contentious debates within both academia and industry on price bubbles in commodity markets and the role of speculators in particular (see, for instance, Gilbert (2008)). Undisputed is the fact that speculative trading is a key feature of modern commodity markets as investors and funds seek to actively and often aggressively search out returns. In this context, this study makes a number of key contributions to the literature.

Quantitative trading opportunities in the crude oil and refined products markets are investigated over the period 2003–2010, with particular focus on WTI, Brent, heating oil and gas oil. A wide range of common- and cross-commodity spreads (including calendar, crack and locational spreads) are considered. Trading strategies are designed so as to exploit any predictability that exists in the unit volumetric spreads. Drawing on the current literature, the novel *optimal* statistical arbitrage trading model of Bertram (2010) is applied for the empirical analysis in this study.* Few papers deal directly with the issue of optimal entry and exit trading signals for statistical arbitrage trading strategies. These include Vidyamurthy (2004), Elliott *et al.* (2005) and Do *et al.* (2006). The approach of Bertram (2010) is very much distinct from these papers. Specifically, modelling a given spread series as a mean-reverting Ornstein–Uhlenbeck process, analytic solutions exist for the optimal entry and exit levels determined through maximising the expected return *per unit time*, which is defined as the ratio of the deterministic return of the strategy to the expected trade cycle time. Statistical arbitrage trading strategies with defined entry and exit levels offer deterministic returns but uncertainty lies in the length of the trade cycles. Normalisation by the expected trade cycle time explicitly accounts for the alternative deterministic returns and stochastic trade cycle times associated with alternative strategies, allowing for consistent cross-comparison.

The full set of commodity data considered leads to 861 spreads in total. The objective of this study is to statistically test the performance of just over 2500 statistical arbitrage style strategies simultaneously, based on the constructed spreads. The hypothesis tests are designed to formally identify those trading strategies that, with statistical significance, outperform a given benchmark in terms of mean daily log-return. The benchmark is defined to have zero mean daily log-return, corresponding to taking no position in a given spread. This introduces the well-known issue of data snooping bias, whereby, under naive analysis, profitable trading strategies may be identified by pure chance alone. This links directly to the broader issue of multiple hypothesis testing in general statistical and econometric applications. Romano *et al.* (2009) provide a detailed exposition of the issues pertaining to multiple hypotheses testing, outlining the key literature in the area. Two recent generalised stepwise techniques proposed in the literature are used to control for data snooping bias. The stepdown procedure of Romano and Wolf (2007) and the balanced stepdown procedure of Romano and Wolf (2010) are applied, both serving as improvements over more conservative single-step approaches, such as the seminal reality check bootstrap test of White (2000) and the superior predictive ability test of Hansen (2005). The generalised procedures offer greater *power*, where power is loosely defined as the ability to reject false null hypotheses. The balanced stepdown procedure offers a further improvement in allowing for balance amongst the hypothesis tests in the sense that each is treated equally in terms of power. Applying these generalised stepwise procedures to control for data snooping bias within this quantitative trading study presents a significant contribution to the literature. Specifically, whereas the conservative single-step procedures

* A number of alternative studies in the area of statistical arbitrage trading have been conducted to date, including Burgess (1999, 2000), Trapletti *et al.* (2002), Vidyamurthy (2004), Whistler (2004), Elliott *et al.* (2005), Gatev *et al.* (2006), Do *et al.* (2006) and Avellaneda and Lee (2010). Andrade *et al.* (2005), Papadakis and Wysocki (2007) and Do and Faff (2009) provide independent verification of the trading rule proposed by Gatev *et al.* (2006). Aldridge (2009), Bowen *et al.* (2010) and Dunis *et al.* (2010) consider high-frequency applications. Other papers of interest include Shleifer and Vishny (1997), Hogan *et al.* (2004) and Lin *et al.* (2006).

allow for statistical inferences around the best-performing trading strategy only, the stepwise procedures recursively identify all of those trading strategies that are profitable with statistical significance.

The reality check bootstrap test of White (2000) has been applied in a number of quantitative trading studies to identify profitability while rigorously controlling for data snooping bias. Sullivan *et al.* (1999), Hsu and Kuan (2005), Qui and Wu (2006), Park and Irwin (2007), and Marshall *et al.* (2008) apply the reality check bootstrap test to evaluate the profitability of a wide range of technical trading rules commonly used in industry. Marshall *et al.* (2008) consider a data set of 15 commodities, including crude oil and heating oil, and conclude that, with the exception of oats, technical trading rules are not profitable. Hsu *et al.* (2009) apply a stepwise extension of the superior predictive ability test of Hansen (2005) to re-evaluate the profitability of technical trading rules. Our study is distinct from these papers in that statistical arbitrage techniques are employed to exploit predictability or mean-reversion in the spreads where it exists.

The study further contributes by allowing for practical comparison of the stepwise and the balanced stepwise procedures in the context of a trading application. The stepwise procedure of Romano and Wolf (2007) improves on single-step procedures with subsequent iterative steps that allow for additional profitable strategies to be identified. The balanced stepwise procedure of Romano and Wolf (2010) likewise improves on single-step procedures with subsequent iterative steps that allow for additional profitable strategies to be identified, but with the added benefit of balance such that trading strategies with large mean daily returns do not dominate those with lower mean daily returns. The balanced stepdown procedure is shown to identify many more profitable strategies than the non-balanced stepdown procedure. However, there exists broad consistency, in the sense that those trading strategies identified as being profitable by the stepdown procedure are also identified in general by the balanced stepdown procedure.

The remainder of the paper is organised as follows. Section 7.2 discusses the crude oil and refined products markets, describing in detail the data used for the empirical analysis. Section 7.3 presents the optimal statistical arbitrage trading model of Bertram (2010). Section 7.4 discusses the issue of data snooping bias and links this to the broader issue of multiple hypothesis testing. The details of the stepdown procedure of Romano and Wolf (2007) and the balanced stepdown procedure of Romano and Wolf (2010) are also given. The empirical analysis conducted in this study is outlined in Section 7.5, defining the formal hypothesis tests and discussing transaction cost and liquidity issues. Section 7.6 presents the results of the empirical analysis and considers robustness to varying transaction costs. Section 7.7 concludes.

7.2 Crude Oil and Refined Products Markets

The importance of crude oil and refined products markets within the modern economy is well appreciated. Increased volatility and uncertainty around prices underlies the active trading of derivatives products to hedge and manage the wide variety of risk exposures of producers and consumers in the energy sector. The trading of futures spreads offer a range of flexible risk management solutions. Calendar spreads involve the trading of futures contracts on the same underlying commodity with different maturities. Such spreads allow producers and suppliers to trade maturity-based price differentials and hedge movements along the futures curve, in particular transitions between backwardation and contango. Crack spreads represent the price differential between a given crude oil and refined product and essentially represent the production margin for refineries. Locational spreads represent the price differential between the same or related commodities trading in different geographical regions. These spreads allow energy market participants to exploit locational arbitrage, subject to transport and associated costs in the case of physical deliveries.

Volatility in the crude oil and refined products markets has attracted much interest in speculative trading opportunities. The quantitative trading of futures spreads in these markets is the primary focus of this study, premised on the existence of predictability and mean-reversion in relative price movements. Under the seminal cost of carry theory (Kaldor 1939), a number of studies examine the term structure relationship between spot and futures prices from both a market efficiency perspective (Crowder

and Hamed 1993, Kellard *et al.* 1999) and a price discovery perspective (Schwarz and Szakmary 1994, Silvapulle and Moosa 1999). Girma and Paulson (1999) and Gjolberg and Johnsen (1999) examine the relation between crude oil and a range of refined products. It is concluded that cointegrating relationships exist, implying predictability in the crack spreads. Kinnear (2002) and Alizadeh and Nomikos (2004) consider locational arbitrage opportunities in the crude oil markets. All of these studies provide motivation for exploring the existence of predictability and mean-reversion in spreads from a trading perspective. Rather than test for these features directly, the approach taken in this study is to apply mean-reversion-based trading strategies to a large data set of spreads and use advanced data snooping control procedures to identify statistically significant profitability.

For the empirical analysis to follow, a comprehensive data set of crude oil and refined product futures contracts are used, comprising WTI and Brent on the crude oil side and gas oil (GO) and heating oil (HO) on the refined products side. These commodities are chosen for this study on the basis of size, importance and liquidity. WTI and HO are both traded on the New York Mercantile Exchange (NYMEX), while Brent and GO are both traded on the Intercontinental Exchange (ICE). The data set covers the 11-year period from 3 January 2000 to 31 December 2010. Most notably, this period covers the record high crude oil prices recorded in 2008 and the subsequent collapse in the latter part of the same year resulting from the global economic crisis, in addition to the gradual recovery in crude oil prices over 2009–2010. All relevant conversions were done to ensure the time series are quoted consistently in dollars per barrel.

The data set includes futures curves for WTI, Brent and HO running from the prompt month up to month 12, with the GO futures curve running from the prompt month up to month 6. These choices are made to ensure sufficient liquidity from a trading perspective, with the construction of the individual time series explicitly taking into account the rolling of futures contracts. Transaction costs and contract liquidity are discussed in more detail later in [Section 7.5](#). For the quantitative spread trading analysis, the full range of common and cross-commodity spreads (including calendar, crack and locational spreads) are considered. With 42 different maturity contracts across the four commodity groups, a total of 861 individual spreads are available for analysis. Finally, non-synchronicity bias is avoided with all time series being observed at the same time of 5.15 p.m. EDT, coinciding with the close of the WTI crude oil market.

7.3 Optimal Statistical Arbitrage Trading Model

This section provides a detailed mathematical exposition of the novel optimal statistical arbitrage trading model of Bertram (2010). The issue of optimal statistical arbitrage trading is approached by first assuming that the spread between two asset log-price series, denoted s_t , is given by the following zero-mean OU process:^{*}

$$ds_t = -\alpha s_t dt + \sigma dW_t, \quad (7.1)$$

with $\alpha, \sigma > 0$ and W_t denoting a Wiener process. Defining the entry and exit levels of the trading strategy by a and m , respectively, a complete trade cycle is the time taken for the spread process to transition from a to m and then return back to a . Formally, the trade cycle time is defined as follows:

$$\mathcal{T} \equiv \mathcal{T}_{a \rightarrow m} + \mathcal{T}_{m \rightarrow a},$$

where $\mathcal{T}_{a \rightarrow m}$ is the time to transition from a to m and $\mathcal{T}_{m \rightarrow a}$ is the time to transition from m to a , and the independence of the two times follows from the Markovian property of the OU process. Given relative transaction costs c , the total log-return from one trade cycle of the statistical arbitrage trading strategy

^{*} The zero-mean assumption does not present any issue in practice. The optimal entry and exit levels obtained can be easily translated to account for a non-zero mean in empirical data.

is given by $r(a, m, c) \equiv m - a - c$, which is deterministic but for which the associated trading cycle time is stochastic. In this context, Bertram (2010) proposes the following expected return per unit time and variance of return per unit time measures:

$$\xi(a, m, c) \equiv \frac{r(a, m, c)}{E(\mathcal{T})},$$

$$\varsigma(a, m, c) \equiv \frac{r^2(a, m, c)V(\mathcal{T})}{E^3(\mathcal{T})},$$

where $E(T) = E(\mathcal{T}_{a \rightarrow m}) + E(\mathcal{T}_{m \rightarrow a})$ is the expected trade cycle time and $V(T) = V(\mathcal{T}_{a \rightarrow m}) + V(\mathcal{T}_{m \rightarrow a})$ is the variance of the trade cycle time. Following a transformation of the OU process to a dimensionless system, and drawing on the first-passage time theory of Thomas (1975), Sato (1977) and Ricciardi and Sato (1988), Bertram (2010) derives the following analytic expressions for $E(T)$, $V(T)$, $\xi(a, m, c)$ and $\varsigma(a, m, c)$:

$$E(\mathcal{T}) = \frac{\pi}{\alpha} \left(\operatorname{Erfi} \left(\frac{m\sqrt{\alpha}}{\sigma} \right) - \operatorname{Erfi} \left(\frac{a\sqrt{\alpha}}{\sigma} \right) \right),$$

$$V(\mathcal{T}) = \frac{w_1 \left(\frac{m\sqrt{2\alpha}}{\sigma} \right) - w_1 \left(\frac{a\sqrt{2\alpha}}{\sigma} \right) - w_2 \left(\frac{m\sqrt{2\alpha}}{\sigma} \right) + w_2 \left(\frac{a\sqrt{2\alpha}}{\sigma} \right)}{\alpha^2},$$

$$\xi(a, m, c) = \frac{\alpha(m - a - c)}{\pi \left(\operatorname{Erfi} \left(\frac{m\sqrt{\alpha}}{\sigma} \right) - \operatorname{Erfi} \left(\frac{a\sqrt{\alpha}}{\sigma} \right) \right)},$$

$$\varsigma(a, m, c) = \alpha(m - a - c)^2 \times \frac{w_1 \left(\frac{m\sqrt{2\alpha}}{\sigma} \right) - w_1 \left(\frac{a\sqrt{2\alpha}}{\sigma} \right) - w_2 \left(\frac{m\sqrt{2\alpha}}{\sigma} \right) + w_2 \left(\frac{a\sqrt{2\alpha}}{\sigma} \right)}{\pi^3 \left(\operatorname{Erfi} \left(\frac{m\sqrt{\alpha}}{\sigma} \right) - \operatorname{Erfi} \left(\frac{a\sqrt{\alpha}}{\sigma} \right) \right)^3}$$

where $\operatorname{Erfi}(z)$ is the imaginary error function,

$$w_1(z) \equiv \left(\frac{1}{2} \sum_{g=1}^{\infty} \Gamma \left(\frac{g}{2} \right) (\sqrt{2z})^g / g! \right)^2$$

$$- \left(\frac{1}{2} \sum_{g=1}^{\infty} (-1)^g \Gamma \left(\frac{g}{2} \right) (\sqrt{2z})^g / g! \right)^2,$$

$$w_2(z) \equiv \sum_{g=1}^{\infty} \Gamma \left(\frac{(2g-1)}{2} \right) \Psi \left(\frac{(2g-1)}{2} \right) (\sqrt{2z})^{(2g-1)} / (2g-1)!$$

and $\Psi(z) \equiv \Psi(z) - \Psi(1)$, with $\Gamma(z)$ and $\Psi(z)$ the gamma and digamma functions, respectively.

With these analytic results in place, it is shown that the optimal entry and exit levels a^* and m^* may be derived by either maximising the expected return per unit time $\xi(a, m, c)$ or the associated per unit time Sharpe ratio. In the former case, it is established that $m^* = -a^*$, with $a^* < 0$ being the root of the equation

$$\exp \left(\frac{\alpha a^2}{\sigma^2} \right) (2a + c) - \sigma \sqrt{\frac{\pi}{\alpha}} \operatorname{Erfi} \left(\frac{a\sqrt{\alpha}}{\sigma} \right) = 0.$$

In the latter case, the risk-free rate of interest r_f is introduced and the associated per unit time Sharpe ratio is given by

$$\begin{aligned}
 S(a, m, c, r_f) &\equiv \frac{\xi(a, m, c) - (r_f/E(\mathcal{T}))}{\sigma} \\
 &= (m - a - c - r_f) \sqrt{\frac{E(\mathcal{T})}{(m - a - c)^2 V(\mathcal{T})}} \\
 &= (m - a - c - r_f) \\
 &\quad \times \sqrt{\frac{\frac{\pi}{\alpha} \left(\operatorname{Erfi}\left(\frac{m\sqrt{\alpha}}{\sigma}\right) - \operatorname{Erfi}\left(\frac{a\sqrt{\alpha}}{\sigma}\right) \right)}{(m - a - c)^2 \frac{w_1 \left(\frac{m\sqrt{2\alpha}}{\sigma}\right) - w_1 \left(\frac{a\sqrt{2\alpha}}{\sigma}\right) - w_2 \left(\frac{m\sqrt{2\alpha}}{\sigma}\right) + w_2 \left(\frac{a\sqrt{2\alpha}}{\sigma}\right)}{\alpha^2}}}.
 \end{aligned}$$

It is established similarly that $m = -a$, $a < 0$, and the optimal entry level a^* follows from maximising the Sharpe ratio expression with this substitution for m .

7.3.1 Remark on Model Specification

The OU process described in Equation (7.1) for the log-price spread series is simplistic in its specification, although it has been used extensively in the finance literature, particularly for interest rate modelling (Vasicek 1977). The OU process assumes a Gaussian distribution for the spread series, with constant volatility, constant long-run mean level, constant speed of mean reversion and symmetry in adjustments from above and from below the long-run mean level. Given the non-Gaussian nature of financial data, the OU process is inherently misspecified. However, the purpose of this study is to dynamically exploit predictability in the spread between the log-prices of related oil and oil-based products. The key advantage of using the OU process to describe this predictability is that it allows for analytic trade execution signals to be determined, which in the broader context provides significant computational efficiencies that are of particular benefit for high-frequency statistical arbitrage trading. As noted by Bertram (2010), the simplicity of the OU process further offers ease in investigating how the important system variables relate to each other, in particular time. Model specification issues for the Bertram (2010) trading model are discussed by Cummins (2011), where it is concluded that the closer a spread series is to normal then the less the model misspecification error will be in general.

7.4 Multiple Hypothesis Testing: Data Snooping Bias

The objective of the study is to formally and statistically test the performance of the optimal statistical arbitrage trading model of Bertram (2010) described in the previous section for the quantitative trading of spreads in the crude oil and refined products markets. This will inevitably involve the testing of a large number of implementations of the trading model simultaneously. This introduces the well-established issue of data snooping bias, whereby, under naive analysis, profitable trading strategies may be identified by pure chance alone. This links directly to the broader issue of multiple hypothesis testing in statistical and econometric applications. The remainder of the section expands on this, and introduces two key generalised techniques that will be used to control for the problem of data snooping bias.

The issue with multiple hypothesis testing is that the probability of false discoveries, i.e. the rejection of true null hypotheses by chance alone, is often significant. There are a number of approaches described in the literature to deal with this multiple comparisons problem and control for the

familywise error rate and variants. Romano *et al.* (2009) provide an excellent summary of the issues and the literature. The familywise error rate (FWER) is defined as the probability that at least one or more false discoveries occur. Consistent with the notation of Romano *et al.* (2009), the following definition is made:

$$FWER_0 = P_0 \{ \text{reject at least one null hypothesis } H_{0,s} : s \in \mathcal{I}(\theta) \},$$

where $H_{0,s}$, $s = 1, \dots, S$, is a set of null hypotheses, and $\mathcal{I}(\theta)$ is the set of true null hypotheses. Controlling the FWER involves setting a significance level $\tilde{\alpha}$ and requiring that $FWER_0 \leq \tilde{\alpha}$. Controlling for the FWER in this way is particularly conservative given that it does not allow even for one false discovery and so is criticised for lacking *power*, where power is loosely defined as the ability to reject false null hypotheses, i.e. identify true discoveries (Romano *et al.* 2009). The greater S , the more difficult it is to make true discoveries.

To deal with this weakness, generalised FWER approaches have been proposed in the literature. The generalised FWER seeks to control for k (where $k \geq 1$) or more false discoveries and, in so doing, allows for greater power in multiple hypothesis testing. The generalised k -FWER is defined as follows:

$$k\text{-FWER}_0 = P_0 \{ \text{reject at least } k \text{ null hypothesis } H_{0,s} : s \in \mathcal{I}(\theta) \}.$$

Towards building a framework to identify profitable trading strategies, with statistical significance, on the set of calendar and crack spreads in this study, the following one-sided hypothesis tests are considered:

$$H_{0,s}, \theta_s \leq 0 \text{ vs. } H_{1,s}, \theta_s > 0.$$

The objective is to control for the multiple comparisons in this scenario through the generalised familywise error rate, which offers greater power while also implicitly accounting for the dependence structure that exists between the tests. This section continues as follows. Section 7.4.1 presents a single-step procedure as described by Romano *et al.* (2009). Section 7.4.2 presents the stepwise procedure of Romano and Wolf (2007), which serves as an improvement on the single-step approach by allowing for subsequent iterative steps to identify additional hypothesis rejections. Finally, Section 7.4.3 presents the balanced stepwise procedure of Romano and Wolf (2010), which is again a marked improvement that allows for balance amongst the hypothesis tests in the sense that each is treated equally in terms of power, i.e. in the identification of true discoveries.

7.4.1 Single-Step Procedure

Assume a set of test statistics $T_{n,s} = \hat{\theta}_{n,s}$ associated with the hypothesis tests, where n is introduced to denote the sample size of the data used for estimation. Letting $A \equiv \{1, \dots, S\}$, the single-step procedure proceeds by rejecting all hypotheses where $T_{n,s} \geq c_{n,A}(1 - \tilde{\alpha}, k)$, and where $c_{n,A}(1 - \tilde{\alpha}, k)$ represents the $(1 - \tilde{\alpha})$ -quantile of the distribution of $k\text{-max}(\hat{\theta}_{n,s} - \theta_s)$ under P_0 . With P_0 unknown, the critical value $c_{n,A}(1 - \tilde{\alpha}, k)$ is also unknown. However, an estimate critical value may be determined using appropriate bootstrapping techniques. That is, the critical value $\hat{c}_{n,A}(1 - \tilde{\alpha}, k)$ is estimated as the $(1 - \tilde{\alpha})$ -quantile of the distribution of $k\text{-max}(\hat{\theta}_{n,s}^* - \hat{\theta}_{n,s})$ for $\hat{P}_0$ an unrestricted estimate of P_0 . See Romano *et al.* (2009) for further technical details.

7.4.2 Stepdown Procedure

The stepwise procedure of Romano and Wolf (2007) improves on the single-step procedure described in the previous section by allowing for subsequent iterative steps to identify additional hypothesis rejections. The stepdown procedure is constructed such that, at each stage, information on the rejected

hypotheses to date is used in re-testing for significance on the remaining hypotheses. Romano and Wolf (2007) describe the following steps to the algorithm.

- **Step 1:** Let A_1 denote the full set of hypothesis indices, i.e. $A_1 \equiv \{1, \dots, S\}$. If the maximum test statistic observed, i.e. $\max(T_{n,s})$, is less than or equal to the estimated critical value $\hat{c}_{n,A_1}(1-\tilde{\alpha}, k)$, then fail to reject all null hypotheses and stop the algorithm. Otherwise, proceed to reject all null hypotheses $H_{0,s}$ for which the associated test statistics exceed the critical value level, i.e. where $T_{n,s} > \hat{c}_{n,A_1}(1-\tilde{\alpha}, k)$.
- **Step 2:** Let R_2 denote the set of indices for the hypotheses rejected in Step 1 and let A_2 denote the indices for those hypotheses not rejected. If the number of elements in R_2 is less than k , i.e. $|R_2| < k$, then stop the algorithm, as the probability of k or more false discoveries is zero in this case. Otherwise, the appropriate critical value to be applied at this stage is calculated as follows:

$$\hat{d}_{n,A_2}(1-\tilde{\alpha}, k) = \max_{I \subseteq R_2, |I|=k-1} \left\{ \hat{c}_{n,K}(1-\tilde{\alpha}, k) : K \equiv A_2 \cup I \right\}.$$

Hence, additional hypotheses from A_2 are rejected if $T_{n,s} > \hat{d}_{n,A_2}(1-\tilde{\alpha}, k)$, $s \in A_2$. If no further rejections are made, then stop the algorithm.

⋮

- **Step j:** Let R_j denote the set of indices for the hypotheses rejected up to Step $(j-1)$ and let A_j denote the indices for those hypotheses not rejected. The critical value to be applied at this stage is calculated as follows:

$$\hat{d}_{n,A_j}(1-\tilde{\alpha}, k) = \max_{I \subseteq R_j, |I|=k-1} \left\{ \hat{c}_{n,K}(1-\tilde{\alpha}, k) : K \equiv A_j \cup I \right\}.$$

Hence, additional hypotheses from A_j are rejected if $T_{n,s} > \hat{d}_{n,A_j}(1-\tilde{\alpha}, k)$, $s \in A_j$. If no further rejections are made, then stop the algorithm.

⋮

From the description of the above algorithm, at each stage j in the stepwise procedure, the hypotheses that are not rejected thus far are re-tested over a smaller population of hypothesis tests than previously. The size of this smaller population is given by $(|A_j| + k - 1)$, which includes all the hypotheses within A_j , in addition to $(k - 1)$ hypotheses drawn from those already rejected, i.e. drawn from R_j . Given that control of the generalised k -FWER is the premise of the procedure, it is expected that there are at most $(k - 1)$ false discoveries amongst the set of hypotheses rejected R_j . However, it is not known which of the rejected hypotheses may represent false discoveries. Hence, it is necessary to circulate through all combinations of R_j , of size $(k - 1)$, in order to obtain the maximum critical value $\hat{d}_{n,A_j}(1-\tilde{\alpha}, k)$ against which to test the hypotheses within A_j . See Romano and Wolf (2007) for further technical details.

7.4.2.1 Operative Method

In requiring to circulate through all subsets of R_j , of size $(k - 1)$, in order to obtain the maximum critical value to apply at each stage of the stepdown procedure, the algorithm can become highly, if not excessively, computationally burdensome. Depending on $|R_j|$ and the value of k , the number of combinations ${}^{|R_j|}C_{k-1}$ can become very large. Romano and Wolf (2007) therefore suggest an operative method that

reduces this computational burden, while at the same time maintaining much of the attractive properties of the algorithm.*

For this, first consider the hypothesis tests rejected up to step $(j - 1)$ and place these in descending order of test statistic, i.e.

$$T_{n,r_1} \geq T_{n,r_2} \geq \dots \geq T_{n,r_{|R_j|}},$$

where $\{r_1, r_2, \dots, r_{|R_j|}\}$ is the appropriate permutation of associated hypothesis test indices that gives this ordering. Now consider a user-defined maximum number of combinations, $N_{\max}$, at each step of the algorithm. Then choose an integer value such that ${}^M C_{k-1} \leq N_{\max}$ and replace the critical value calculation at each step j of the algorithm with the following:

$$\hat{d}_{n,A_j}(1-\tilde{\alpha}, k) = \max_{I \subseteq \{r_{\max(1, |R_j|-M+1)}, \dots, r_{|R_j|}\}, |I|=k-1} \{\hat{c}_{n,K}(1-\tilde{\alpha}, k) : K \equiv A_j \cup I\}.$$

What this serves to do is to replace circulating through all the hypothesis tests rejected to date with that of circulating through only the M least-significant hypothesis tests rejected. Of course, in the case where $M \geq |R_j|$, then this amounts to circulating through all the hypotheses rejected. Although this approach is premised on the assumption that the (up to $k - 1$) false discoveries lie within the least-significant hypotheses rejected so far, it does offer significant computational efficiencies for the algorithm. It is this operative method that is used for the empirical analysis in subsequent sections.

7.4.3 Balanced Stepdown Procedure

Whereas the stepwise procedure of the previous section is an improvement on the single-step procedure of [Section 7.4.1](#), it does not offer by construction balance in the sense that each hypothesis test is treated equally in terms of power. The balanced stepwise procedure of Romano and Wolf (2010) addresses this issue.

Introducing some notation, let $H_{n,s}(\cdot, P_\theta)$ denote the distribution function of $(\hat{\theta}_{n,s} - \theta_s)$ and let $c_{n,s}(\tilde{\gamma})$ denote the $\tilde{\gamma}$ -quantile of this distribution. The confidence interval

$$\{\theta_s : \hat{\theta}_{n,s} - \theta_s \leq c_{n,s}(\tilde{\gamma})\}$$

then has coverage probability $\tilde{\gamma}$. Balance is the property that the marginal confidence intervals for a population of S simultaneous hypothesis tests have the same probability coverage. Within the context of controlling the generalised k -FWER, the overall objective is to ensure that the simultaneous confidence interval covers all parameters $\theta_s, s = 1, \dots, S$, except for at most $(k - 1)$ of them, for a given limiting probability $(1 - \tilde{\alpha})$, while at the same time ensuring balance (at least asymptotically). So, what is sought is that

$$\begin{aligned} & P_\theta \left\{ \hat{\theta}_{n,s} - \theta_s \leq c_{n,s}(\tilde{\gamma}) \text{ for all but at most } (k-1) \text{ of the hypotheses} \right\} \\ & \equiv P_\theta \left\{ H_{n,s}(\hat{\theta}_{n,s} - \theta_s, P_\theta) \leq \tilde{\gamma} \text{ for all but at most } (k-1) \text{ of the hypotheses} \right\} \\ & \equiv P_\theta \left\{ k - \max \left(H_{n,s}(\hat{\theta}_{n,s} - \theta_s, P_\theta) \right) \leq \tilde{\gamma} \right\} = 1 - \tilde{\alpha}. \end{aligned}$$

Letting $L_{n\{1, \dots, S\}}(k, P_\theta)$ denote the distribution of $k - \max \left(H_{n,s}(\hat{\theta}_{n,s} - \theta_s, P_\theta) \right)$, the appropriate choice of the coverage probability $\tilde{\gamma}$ is then $L_{n\{1, \dots, S\}}^{-1}(1 - \tilde{\alpha}, k, P_\theta)$.

* The generic algorithm offers a number of attractive features. Firstly, the generic algorithm is conservative in its rejection of hypotheses. Secondly, the generic algorithm also allows for finite sample control of the k -FWER under P_θ . And thirdly, the bootstrap construction is such that the generic algorithm provides asymptotic control in the case of contiguous alternatives. Romano and Wolf (2007) provide a more detailed discussion.

As before, given that P_0 is unknown, it is necessary to use appropriate bootstrapping techniques to generate an estimate of the coverage probability $L_{n,\{1,\dots,S\}}^{-1}(1-\tilde{\alpha}, k, \hat{P}_0)$, under $\hat{P}_0$. Therefore, from this development it is possible to define the simultaneous confidence interval

$$\left\{ \theta_s : \hat{\theta}_{n,s} - \theta_s \leq H_{n,s}^{-1} \left(L_{n,\{1,\dots,S\}}^{-1}(1-\tilde{\alpha}, k, \hat{P}_0), \hat{P}_0 \right) \right\}.$$

The right-hand side of the above inequality will form the basis of the critical value definitions used within the stepdown procedure. See Romano and Wolf (2010) for further technical details. Note that although the above development was made assuming the full set of hypothesis tests, it equally applies to any subset $K \subseteq \{1, \dots, S\}$. Hence, the balanced stepwise algorithm may now be described as follows.

- **Step 1:** Let A_1 denote the full set of hypothesis indices, i.e. $A_1 \equiv \{1, \dots, S\}$. If, for each hypothesis test, the associated test statistic $T_{n,s}$ is less than or equal to the corresponding critical value estimate $\hat{c}_{n,A_1,s}(1-\tilde{\alpha}, k) \equiv H_{n,s}^{-1} \left(L_{n,A_1}^{-1}(1-\tilde{\alpha}, k, \hat{P}_0), \hat{P}_0 \right)$, then fail to reject all null hypotheses and stop the algorithm. Otherwise, proceed to reject all null hypotheses $H_{0,s}$ for which the associated test statistics exceed the critical value level, i.e. where $T_{n,s} > \hat{c}_{n,A_1,s}(1-\tilde{\alpha}, k)$.
- **Step 2:** Let R_2 denote the set of indices for the hypotheses rejected in Step 1 and let A_2 denote the indices for those hypotheses not rejected. If the number of elements in R_2 is less than k , i.e. $|R_2| < k$, then stop the algorithm, as the probability of k or more false discoveries is zero in this case. Otherwise, the appropriate critical value to be applied for each hypothesis test s at this stage is calculated as follows:

$$\hat{d}_{n,A_2,s}(1-\tilde{\alpha}, k) = \max_{I \subseteq R_2, |I|=k-1} \left\{ \hat{c}_{n,K,s}(1-\tilde{\alpha}, k) : K \equiv A_2 \cup I \right\}.$$

Hence, additional hypotheses from A_2 are rejected if $T_{n,s} > \hat{d}_{n,A_2,s}(1-\tilde{\alpha}, k)$, $s \in A_2$. If no further rejections are made, then stop the algorithm.

⋮

- **Step j:** Let R_j denote the set of indices for the hypotheses rejected up to Step $(j-1)$ and let A_j denote the indices for those hypotheses not rejected. The appropriate critical value to be applied for each hypothesis test s at this stage is calculated as follows:

$$\hat{d}_{n,A_j,s}(1-\tilde{\alpha}, k) = \max_{I \subseteq R_j, |I|=k-1} \left\{ \hat{c}_{n,K,s}(1-\tilde{\alpha}, k) : K \equiv A_j \cup I \right\}.$$

Hence, additional hypotheses from A_j are rejected if $T_{n,s} > \hat{d}_{n,A_j,s}(1-\tilde{\alpha}, k)$, $s \in A_j$. If no further rejections are made, then stop the algorithm.

⋮

Similar to the stepwise algorithm of the previous section, at each stage j in the stepwise procedure, the hypotheses that are not rejected thus far are re-tested over a smaller population of hypothesis tests than previously. The size of this smaller population is given by $(|A_j| + k - 1)$, which includes all the hypotheses within A_j , in addition to $(k - 1)$ hypotheses drawn from those hypotheses already rejected, i.e. drawn from R_j . Given that control of the generalised k -FWER is the premise of the procedure, it is expected that there are at most $(k - 1)$ false discoveries amongst the set of hypotheses rejected R_j . However, it is not known which of the rejected hypotheses may represent false discoveries. Hence, it is necessary to

circulate through all combinations of R_p of size $(k - 1)$, in order to obtain the appropriate critical values. Where the algorithm departs significantly from the previous section is that a maximum critical value $\hat{d}_{n,A_j,s}(1 - \tilde{\alpha}, k)$ must be determined for each hypothesis test s . This adds an additional layer of computational burden on the algorithm.

7.4.3.1 Operative Method

Similar to the stepdown procedure of [Section 7.4.2](#), the need to circulate through all subsets of R_p of size $(k - 1)$, in order to obtain, in this case, a set of maximum critical values to apply at each stage of the stepdown procedure, means the algorithm can become excessively computationally burdensome. Romano and Wolf (2010) therefore suggest an operative method that reduces this computational burden in the spirit of that proposed by the authors for the stepdown procedure (Romano and Wolf 2007).

It is first necessary to be able to order the hypothesis tests rejected up to step $(j - 1)$ in terms of significance. To this end, it is noted that marginal p -values can be obtained as follows:

$$\hat{p}_{n,s} \equiv 1 - H_{n,s}(\hat{\theta}_{n,s}, \hat{P}_0).$$

This gives the following ascending order for the significance of the hypothesis tests:

$$\hat{p}_{n,r_1} \leq \hat{p}_{n,r_2} \leq \dots \leq \hat{p}_{n,r_{|R_j|}},$$

where $\{r_1, r_2, \dots, r_{|R_j|}\}$ is the appropriate permutation of associated hypothesis test indices that gives this ordering. As before, a maximum number of combinations, $N_{\max}$, at each step of the algorithm is defined. Then an integer value M is chosen such that ${}^M C_{k-1} \leq N_{\max}$, leading to the calculation of the critical values as follows:

$$\hat{d}_{n,A_j,s}(1 - \tilde{\alpha}, k) = \max_{I \subseteq \{r_{\max}(1, |R_j| - M + 1), \dots, r_{|R_j|}\}, |I| = k-1} \{\hat{c}_{n,K,s}(1 - \tilde{\alpha}, k) : K \equiv A_j \cup I\}.$$

The rationale for this approach is exactly the same as that described in [Section 7.4.2.1](#), with the same effect being the introduction of significant computational efficiencies to the algorithm. It is this operative method that is used for the empirical analysis in subsequent sections.

7.5 Outline of Empirical Analysis

With the objective being to formally test the performance of the optimal statistical arbitrage trading model for the quantitative trading of spreads in the crude oil and refined products markets, data snooping bias presents a real issue that needs to be controlled. Furthermore, the performance tests are correlated by nature of the pairings and the markets considered, and so any testing procedure needs to be sufficiently flexible to account for this dependence structure. The stepdown and balanced step-down procedures described in the previous section offer a more generalised and flexible approach to controlling data snooping bias than the commonly used reality check bootstrap test of White (2000) or the superior predictive ability test of Hansen (2005). As identified by Romano *et al.* (2009), both tests lack power in the sense that they control only the familywise error rate (i.e. $k = 1$), and are further restrictive by being only single-step procedures. Additionally, the reality check bootstrap test of White (2000) does not consider studentised test statistics and so by construction suffers from lack of balance. The stepdown and balanced stepdown procedures, in contrast, control the generalised familywise error rate using stepwise procedures, with the latter also offering balance by construction. Given that the test statistics used in this study are not studentised, the balanced stepdown procedure is considered superior to the stepdown procedure as it is asymptotically equivalent to studentisation. Hence, as each test is treated

equally, from a practitioner perspective the balanced stepdown procedure shows no bias between the trading strategies considered. This novel application of the stepdown and balanced stepdown procedures to the problem of quantitative trading within the energy markets represents a key contribution of this study to the literature.

In implementing the stepdown and balanced stepdown procedures, it is first necessary to identify the formal hypothesis tests that are to be applied. The approach taken in this study is to formally identify those spread trading strategies that, with statistical significance, outperform a given benchmark in terms of mean daily log-return. The benchmark is defined as in Hsu and Kuan (2005) as equivalent to taking no position in the spread, and so the hypothesis tests look to identify departures from the zero mean daily log-return. Letting ζ_s denote the daily log-return of trading strategy s and $\theta_s \equiv E(\zeta_s)$, then the hypothesis tests may be formalised as follows:

$$H_{0,s} : \theta_s \leq 0 \text{ vs. } H_{1,s} : \theta_s > 0,$$

for a full set of hypothesis tests $\{1, \dots, S\}$, where S is set equal the number of tests considered. The appropriate estimate of θ_s is the mean daily log-return observed on trading strategy s over a given historical period with n daily observations. Letting $\zeta_{t,s}$, $t = 1, \dots, n$, denote the daily log-return of trading strategy s at time t , the estimate $\hat{\theta}_{n,s} = \sum_{t=1}^n \zeta_{t,s} / n$. Given that θ_s is unknown, implementing the stepdown and balanced stepdown procedures requires use of appropriate bootstrapping techniques. This involves replacing the true specification $(\hat{\theta}_{n,s} - \theta_s)$ with the estimates $(\hat{\theta}_{n,s}^*(b) - \hat{\theta}_{n,s})$, where $\hat{\theta}_{n,s}^*(b)$ are bootstrap estimates of θ_s and $b = 1, \dots, B$ are the indices for the bootstrap samples.

With the hypothesis tests formalised, the next stage is to implement the optimal statistical arbitrage trading model on each spread within the data set. For this it is necessary to fit the following general OU process to each spread time series:

$$ds_t = \alpha(\mu - s_t)dt + \sigma dW_t,$$

in order to obtain parameter estimates $\hat{\alpha}$, $\hat{\mu}$ and $\hat{\sigma}$ for generating the entry and exit trading signals as described in [Section 7.3](#). Adjusting the signals from the trading model for the non-zero long-run mean level $\hat{\mu}$ is straightforward, giving the effective entry and exit levels $\hat{a} \equiv a + \hat{\mu}$ and $\hat{m} \equiv m + \hat{\mu}$, respectively, where by definition of a and m , $\hat{a} < \hat{m}$. To optimise trading performance and exploit the mean-reversion fully (where it exists in the data), trades will be executed to exploit the transition from $\hat{a} \rightarrow \hat{m}$ and the transition back from $\hat{m} \rightarrow \hat{a}$. Therefore, $\hat{a}$ and $\hat{m}$ will herein be referred to merely as trading signals, alternating interpretation between entry and exit signals. The trading strategies are constructed to go long (short) the spread when the trading signal $\hat{a}(\hat{m})$ is breached and then close out the position once $\hat{m}(\hat{a})$ is breached. In this context, it is important to emphasise that the results reported in the next section reflect the aggregation of these *two* trading approaches.

A number of alternative in-sample and out-of-sample periods are considered for estimation and evaluation of the spread trading strategies. Specifically, eight separate one-year out-of-sample periods are considered for evaluation of trading strategy performance, namely, each year over the period 2003–2010. These choices allow for a wider testing of the optimal statistical arbitrage trading model and an examination of trading profitability over time. Three separate in-sample periods are then considered for the estimation of the trading models and the generation of trading signals to be applied out-of-sample. One-year, two-year and three-year periods are considered for estimation, where, for convenience of construction, 252 trade days are assumed in each year. So, for example, the two-year in-sample period includes the 2×252 trade dates prior to the start of the out-of-sample period. Given the daily frequency of the data these choices are seen as reasonable to capture consistent mean-reversion effects, while at the same time examining the impact of alternative estimation periods. Hence, for each out-of-sample period, a total of 2583 (i.e. 3×861) individual trading strategies are tested simultaneously.

To complete the setup of the empirical analysis, it is finally necessary to discuss the choice of generalising parameter k and the probability parameter $\tilde{\alpha}$ to be used within the stepwise and balanced stepwise procedures. To ensure tight control of the number of false discoveries, while at the same time offering power to the tests, k is chosen to ensure that no more than 1% of the tests considered represent false discoveries. Hence, based on a population of $S = 2583$ tests, k is set equal to $\lceil S \times 1\% \rceil = 26$. The significance level $\tilde{\alpha}$ chosen is 5%, such that the implementation of the stepwise and balanced stepwise procedures amount to ensuring that

$$P_{\theta} \{k = 26 \text{ or more false discoveries}\} \leq \tilde{\alpha} = 5\%.$$

7.5.1 Remark on Transaction Costs

Transaction costs are considered and incorporated into the analysis via the transaction costs parameter c detailed in [Section 7.3](#). Given the unavailability to the authors of detailed bid–ask spread information on the underlying futures contracts over the full sample period, the authors draw on the literature to estimate transaction costs. Laws *et al.* (2008) apply filter techniques for the trading of spreads comprised from WTI, Brent and HO using data over the period 1995–2005. The relative transactions costs reported by the authors for these three commodities are presented in Table 7.1. These transaction costs represent the average of four intraday spreads. Without information on bid–ask spreads in the GO market, the relative transaction costs of the HO market are assumed to be a reasonable proxy. Hence, for each spread transaction, the total relative transactions costs are given by the sum of the costs for the individual legs.

The transaction costs are assumed to apply across the full futures curve, which eases the analysis but is clearly restrictive in the sense that transaction costs are likely to increase with increasing maturity. Additionally, given that the sample period of Laws *et al.* (2008) does not coincide with this study, the transactions costs assumed may not be representative. In light of these limitations, the robustness of the empirical results to varying the transaction costs parameter c is examined at the close of the next section. In particular, a range of higher transaction cost levels are assumed and the analysis repeated in order to determine the impact on average daily returns and the number of profitable trades identified.

7.5.2 Remark on Liquidity

An additional issue of practitioner relevance for this study is the level of liquidity across the various futures curves. The ability to execute large volumes of trades based on quantitative trading signals is a key consideration for industry. With only volume data available to the authors, volume-based measures are used here to provide insights into average liquidity across the sample data. Two common measures are calculated for each commodity and maturity futures contract, broken down by each of the 2003–2010 trading years. These measures are the *five-day turnover rate* and the *five-day Hui-Heubel liquidity ratio*. A more comprehensive study of liquidity would be extremely beneficial but, due to data constraints, is beyond the scope of this paper. The issue of liquidity is not considered a source of bias for the stylised trading analysis of this paper given the unit volumetric assumption made in the spread trading design.

TABLE 7.1 Estimated Relative Transactions Costs

WTI	0.0940%
Brent	0.0289%
Heating oil	0.2482%
Gas oil	0.2482%

Sarr and Lybek (2002) provide a detailed overview of liquidity measures, including the two volume-based measures considered here. The turnover rate captures the number of times the outstanding volume changes hands over a specified period. The five-day turnover rate (TR) is used here to give a measure of liquidity over five consecutive trade dates and is defined formally as follows:

$$TR = V / (S \times \bar{P}),$$

where S is the outstanding volume, $\bar{P}$ is the average closing price over the five days and V is the monetary volume value, given by

$$V = \sum_i P_i \times Q_i,$$

where P_i and Q_i are the closing prices and trade volume on date i . The greater the turnover rate, then the greater the depth to the market.

The Hui-Heibel liquidity ratio, on the other hand, seeks to incorporate the impact of trading on price. It considers the percentage price differential between the highest and lowest prices achieved over a specified period relative to the monetary volume value. The five-day Hui-Heibel ratio is used and is defined formally as follows:

$$L_{hh} \equiv [(P_{\max} - P_{\min}) / P_{\min}] / TR,$$

where $P_{\max}$ and $P_{\min}$ are the maximum and minimum closing prices recorded over the five trade days and TR is the five-day turnover rate already defined. The lower the value of the Hui-Heibel liquidity ratio, then the greater the level of liquidity.

Tables A.1–A.4 present, for each commodity and each trading year, the average values of the calculated turnover and Hui-Heibel measures. For the most part, all four commodities show declining liquidity across the futures curve, with high liquidity over the short-dated contracts in particular. Despite the differences in the size of the various markets, there are relatively comparable levels of liquidity across the futures curves based on both measures. The turnover rates also appear to suggest increasing levels of liquidity over the eight-year period, with the Hui-Heibel liquidity ratio results showing broad agreement with this observation.

7.6 Empirical Results

Performing the empirical analysis as set out in the previous section leads to a number of interesting observations. Tables 7.2 and 7.3 present trading performance results based on applying the stepdown and balanced stepdown procedures respectively to control for data snooping bias. For each trading year, average daily returns, Sharpe ratios and trade lengths are reported for three specific categories, namely the top 10, top 20 and all trading strategies identified by the procedures as being profitable with statistical significance. Also reported for each category are the average estimates for the structural parameters of the underlying OU process. The final column in each table gives the total number of profitable trading strategies identified in each year.

It is again important to emphasise that the results reported reflect the aggregation of *two* trading approaches, one taking long positions in the spreads to exploit upward movements between the trading signals and the second taking short positions to exploit downward movements between the trading signals. Many profitable trading strategies are identified, which in many instances report Sharpe ratios that exceed 2 and in some instances are even close to 4. For the top 10 and top 20 categories, average daily returns fall within the approximate range of 0.07–0.35% for most years, with 2009 showing exceptional

TABLE 7.2 Average Empirical Results: Stepdown Procedure

Year		Avg. daily ret.	Avg. Sharpe ratio	Avg. trade length (days)	Avg. α	Avg. μ	Avg. σ	# Profitable strategies
2003	Top 10	0.00361	2.94	15.09	59.367	-0.094	0.259	15
	Top 20	0.00345	2.92	15.81	53.443	-0.103	0.252	
	All	0.00345	2.92	15.81	53.443	-0.103	0.252	
2004	Top 10	0.00299	2.44	18.74	50.529	-0.098	0.295	28
	Top 20	0.00275	2.30	18.88	52.578	-0.096	0.279	
	All	0.00265	2.24	19.58	54.128	-0.099	0.275	
2005	Top 10	0.00259	2.33	19.08	26.167	-0.169	0.275	14
	Top 20	0.00253	2.22	19.81	23.461	-0.172	0.280	
	All	0.00253	2.22	19.81	23.461	-0.172	0.280	
2006	Top 10	0.00241	2.41	19.36	34.984	-0.078	0.287	35
	Top 20	0.00220	2.19	21.28	34.765	-0.092	0.289	
	All	0.00195	1.98	24.22	33.899	-0.102	0.279	
2007	Top 10	0.00293	2.93	12.62	83.288	-0.127	0.233	54
	Top 20	0.00272	2.75	13.33	87.855	-0.120	0.226	
	All	0.00233	2.32	16.22	79.593	-0.121	0.227	
2008	Top 10	-	-	-	-	-	-	-
	Top 20	-	-	-	-	-	-	-
	All	-	-	-	-	-	-	-
2009	Top 10	0.00547	3.95	9.13	92.993	0.010	0.260	-
	Top 20	0.00499	3.66	10.05	86.505	0.007	0.256	-
	All	0.00456	3.32	11.14	95.267	0.008	0.258	31
2010	Top 10	-	-	-	-	-	-	-
-	Top 20	-	-	-	-	-	-	-
-	All	-	-	-	-	-	-	-

results in the range of 0.5–0.55% driven by the locational HO–GO spreads. The associated trade lengths lie in the approximate range of 9–55 trade dates, with the shortest horizons being reported for 2009. The years 2003 and 2009 prove to be particularly successful relative to other years with Sharpe ratios close to or in excess of 3. For the years 2008 and 2010, the stepdown procedure actually fails to identify any profitable strategies, with the balanced stepdown procedure only extracting eight and 23 profitable strategies, respectively. The numbers are much lower than in other years. Average daily returns in excess of 0.2% are reported for 2008 but the associated Sharpe ratios are seen to fall below 2 on average. The discussion to follow will explore further these observations for 2008 in the context of the differences between the stepdown and balanced stepdown procedures. Similar poor performance is observed for the year 2010, where average daily returns can be seen to be at the lower range of approximately 0.07%, with Sharpe ratios again below 2 on average.

Substantial differences are observed between the stepdown and balanced stepdown procedures in terms of the total number of profitable trading strategies identified, despite showing broad consistency within the top 10 and top 20 categories. The balanced stepdown procedure manages to identify between 1.5 and 9.5 times the number of profitable strategies and succeeds in identifying strategies in 2008 and 2010 where the stepdown procedure fails. This observation reflects the manner in which the balanced stepdown procedure treats each hypothesis test equally in terms of power, whereas the stepdown procedure is biased towards those trading strategies with larger average daily returns. The dramatic reduction in the number of profitable strategies in 2008 relative to 2007 and the years previous reflects the collapse in crude oil and other commodity prices resulting from the credit crisis

TABLE 7.3 Average Empirical Results: Balanced Procedure

Year		Avg. daily ret.	Avg. Sharpe ratio	Avg. trade length (days)	Avg. α	Avg. μ	Avg. σ	# Profitable strategies
2003	Top 10	0.00354	2.96	14.28	65.130	-0.089	0.263	142
	Top 20	0.00314	2.87	14.84	66.133	-0.105	0.253	
	All	0.00164	2.16	27.24	28.997	-0.078	0.183	
2004	Top 10	0.00298	2.47	17.58	53.573	-0.096	0.281	52
	Top 20	0.00272	2.34	20.09	47.738	-0.103	0.279	
	All	0.00174	1.84	28.21	33.896	-0.068	0.197	
2005	Top 10	0.00258	2.33	19.46	25.491	-0.167	0.273	71
	Top 20	0.00246	2.22	20.71	23.742	-0.160	0.276	
	All	0.00202	2.05	23.71	23.444	-0.170	0.251	
2006	Top 10	0.00241	2.41	19.36	34.984	-0.078	0.287	46
	Top 20	0.00220	2.19	21.28	34.765	-0.092	0.289	
	All	0.00164	2.03	24.93	36.322	-0.093	0.237	
2007	Top 10	0.00293	2.93	12.62	83.288	-0.127	0.233	82
	Top 20	0.00272	2.75	13.33	87.855	-0.120	0.226	
	All	0.00199	2.18	23.83	66.403	-0.120	0.219	
2008	Top 10	0.00244	1.86	18.26	99.863	0.001	0.232	8
	Top 20	0.00244	1.86	18.26	99.863	0.001	0.232	
	All	0.00244	1.86	18.26	99.863	0.001	0.232	
2009	Top 10	0.00547	3.95	9.13	92.993	0.010	0.260	63
	Top 20	0.00499	3.66	10.05	86.505	0.007	0.256	
	All	0.00264	2.43	24.38	46.336	0.011	0.174	
2010	Top 10	0.00092	1.91	43.88	55.796	-0.014	0.262	23
	Top 20	0.00066	1.82	55.43	30.522	-0.004	0.184	
	All	0.00060	1.73	56.54	26.737	-0.005	0.164	

shock and concerns over global commodity demand. The effect of the collapse in commodity prices manifests as a structural shift in the range of the spreads over 2008 relative to 2007, 2006 and 2005. Therefore, large divergences from the long-run mean levels estimated in-sample are observed for the majority of spreads. These divergences lead to significant losses for the majority of trading strategies out-of-sample in 2008.

For further insights, Appendix B presents the top 10 tradable spreads identified each year and gives the associated average daily return, Sharpe ratio and trade length results. The simplicity of the OU process means parameter estimation may be performed using OLS linear regression. In addition to parameter estimates, *F*-test *p*-values are provided to give an indication of goodness-of-fit of the model. Crack spreads and HO-GO locational spreads can be seen to dominate the composition of the top 10 strategies. The crack spreads include the common-exchange WTI-HO and Brent-GO cracks in addition to the cross-exchange WTI-GO crack. For the most part, the spreads involve either common, adjacent or at least near-by dated contracts. Interestingly, not a single calendar spread makes up the top 10 strategies for any of the years, while locational Brent-WTI spreads only appear in the mix for 2010 (with the exception of one observation in 2003). In terms of the in-sample estimation period used to generate the trading signals, the top 10 profitable strategies do not show any particular pattern. The one-, two- and three-year in-sample estimation periods are all represented within the top 10 strategy groups. Across the eight years, the one-year in-sample period does marginally come out on top, representing approximately 40% of the profitable strategies, with more or less an even split between the two- and three-year estimation periods.

7.6.1 Robustness to Transaction Costs

As outlined in [Section 7.5](#), the levels of transaction costs assumed for the analysis thus far have been taken from Laws *et al.* (2008), with the HO transaction costs assumed to apply for the GO market as well. Given that the sample period used by these authors does not coincide with this study, the transactions costs assumed may not be representative. Therefore, this section now examines the robustness of trading strategy performance to varying transaction cost levels. The transaction parameter c used in the trading model is allowed to range between lower and higher levels than those outlined in [Table 7.1](#), with the objective being to determine the impact on average daily returns and the number of profitable trades identified. For ease of analysis and exposition, a set of transaction cost factors is considered. Each factor is applied to the transaction costs in [Table 7.1](#) and the empirical analysis of the previous section is then repeated. The set of transaction cost factors assumed is $\{0.8, 0.9, \dots, 1.9, 2\}$, allowing transaction costs to vary between 80% and 200% of those assumed so far.

[Figures 7.1–7.4](#) present plots of the average daily returns and the number of profitable trading strategies for each out-of-sample trading year considered. Broadly speaking, it can be seen that the returns and the number of strategies decline with increasing transaction cost levels. However, these declines are not monotonic, reflecting that the calculation of the trading signals is a function of the transaction cost parameter c as outlined in [Section 7.3](#). So as the transaction cost parameter c varies, then so too does the positioning of the signals. Between the lowest and highest transaction cost factors (i.e. 0.8 and 2, respectively), average daily returns in annualised terms decline by approximately 10–15% for 2004, 2006, 2007 and 2009 and by 4–7% for 2003, 2005 and 2010. In the case of trading year 2008, it has been seen already that the number of profitable strategies collapses relative to other years, with only eight trading strategies identified as profitable from the full data set. For the highest transaction cost factor this reduces to a single strategy. The decline in average daily returns in annualised terms between the lowest and highest transaction cost factors is a significant 44%.

7.7 Conclusion

This study examines the quantitative trading of spreads in the crude oil (WTI and Brent) and refined products (heating oil and gas oil) markets, making a number of contributions to the literature. Firstly, the novel statistical arbitrage trading model of Bertram (2010) is applied to a wide range of spreads (including calendar, crack and locational spreads), representing a comprehensive empirical analysis of this model. The model leads to profitable spread trading and it is shown that performance is quite robust to varying transaction costs. Secondly, generalised stepwise procedures are used to control for data snooping bias within the quantitative trading application. The stepdown procedure of Romano and Wolf (2007) and the balanced stepdown procedure of Romano and Wolf (2010) are applied, both serving as improvements over more conservative single-step approaches, such as the reality check bootstrap test of White (2000) and the superior predictive ability test of Hansen (2005). The generalised procedures offer greater power to reject false null hypotheses, with the balanced stepdown procedure offering equal treatment in the identification of profitable strategies. Profitable trading strategies are identified, with results reflecting the aggregation of taking long and short positions in the spreads. For the top 10 and top 20 categories, average daily returns fall within the approximate range of 0.07–0.55%, with trade lengths of 9–55 days and Sharpe ratios of between 2 and 4 in many cases. Thirdly, the study allows for practical comparison of the stepwise and the balanced stepwise procedures in the context of a trading application. The balanced stepdown procedure is unbiased in its approach and is shown to identify many more profitable trading strategies compared with the non-balanced stepdown procedure. For instance, a collapse in the number of profitable trading strategies is seen in 2008, reflecting the impact of the credit crisis and the distortion of spreads relative to previous years. Whereas the stepdown procedure fails to identify any profitable strategies, the balanced procedure is successful in doing so.




FIGURE 7.1 Transaction cost robustness: 2003 and 2004.




FIGURE 7.2 Transaction cost robustness: 2005 and 2006.




FIGURE 7.3 Transaction cost robustness: 2007 and 2008.




FIGURE 7.4 Transaction cost robustness: 2009 and 2010.

References

- Aldridge, I., *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, 2009 (Wiley: Hoboken, NJ).
- Alizadeh, A.H. and Nomikos, N.K., Cost of carry, causality and arbitrage between oil futures and tanker freight markets. *Transport. Res. Part E*, 2004, **40**, 297–316.
- Andrade, S., di Pietro, V. and Seasholes, M., Understanding the profitability of pairs trading. Working Paper, UC Berkley, 2005.
- Avellaneda, M. and Lee, J.H., Statistical arbitrage in the US equities market. *Quant. Finance*, 2010, **10**, 761–782.
- Bertram, W.K., Analytic solutions for optimal statistical arbitrage trading. *Physica A*, 2010, **389**, 2234–2243.
- Bowen, D., Hutchinson, M.C. and O'Sullivan, N., High-frequency equity pairs trading: Transaction costs, speed of execution, and patterns in returns. *J. Trading*, 2010, **5**, 1–38.
- Burgess, A.N., A computational methodology for modelling the dynamics of statistical arbitrage. PhD Thesis, London Business School, 1999.
- Burgess, A.N., Statistical arbitrage models on the FTSE 100. In *Computational Finance*, edited by Y. Abu-Mustafa, B. LeBaron, A.W. Lo and A.S. Weigend, 2000 (MIT Press: Cambridge, MA).
- Crowder, W. and Hamed, A., A cointegration test for oil futures market efficiency. *J. Fut. Mkts*, 1993, **13**, 933–941.
- Cummins, M., Optimal statistical arbitrage: A model specification analysis on ISEQ equity data. *Irish Account. Rev.*, 2011, **17**, 21–40.
- Do, B. and Faff, R., Does naïve pairs trading still work? Working Paper, Monash University, 2009.
- Do, B., Faff, R. and Hamza, K., A new approach to modeling and estimation for pairs trading. Working Paper, Monash University, 2006.
- Dunis, C.L., Giorgini, G., Laws, J. and Rudy, J., Statistical arbitrage and high-frequency data with an application to Eurostoxx 50 equities. Working Paper, Liverpool Business School, 2010.
- Elliott, R.J., Van der Hoek, J. and Malcolm, W.P., Pairs trading. *Quant. Finance*, 2005, **5**, 271–276.
- Gatev, E., Goetzmann, W.N. and Rouwenhorst, K.G., Pairs trading: Performance of a relative-value arbitrage rule. *Rev. Financ. Stud.*, 2006, **19**, 797–827.
- Gilbert, C.L., Commodity speculation and commodity investment. Working Paper, Universita Degli Studi di Trento, 2008.
- Girma, P.B. and Paulson, A.S., Risk arbitrage opportunities in petroleum futures spreads. *J. Fut. Mkts*, 1999, **19**, 931–955.
- Gjolberg, O. and Johnsen, T., Risk management in the oil industry: Can information on long-run equilibrium prices be utilized? *Energy Econ.*, 1999, **21**, 517–527.
- Gregoriou, G.N., Huber, G. and Kooli, M., Performance and persistence of commodity trading advisors: Further evidence. *J. Fut. Mkts*, 2009, **30**, 725–752.
- Hansen, P.R., A test for superior predictive ability. *J. Bus. Econ. Statist.*, 2005, **23**, 365–380.
- Hogan, S., Jarrow, R., Teo, R. and Warachka, M., Testing marking efficiency using statistical arbitrage with applications to momentum and value strategies. *J. Financ. Econ.*, 2004, **73**, 525–565.
- Hsu, P.H. and Kuan, C.M., Re-examining the profitability of technical analysis with White's reality check. Working Paper, 2005.
- Hsu, P.H., Hsu, Y.C. and Kuan, C.M., Testing the predictive ability of technical analysis using a new step-wise test without data snooping bias. Working Paper, 2009.
- Kaldor, N., Speculation and economic stability. *Rev. Econ. Stud.*, 1939, **7**, 1–27.
- Kanamura, T., Rachev, S.T. and Fabozzi, F.J., A profit model for spread trading with an application to energy futures. *J. Trading*, 2010, **5**, 48–62.
- Kavussanos, M.G. and Alizadeh, A., The expectations hypothesis of the term structure and risk premia in dry bulk shipping freight markets. *J. Transp. Econ. Policy*, 2002, **36**, 267–304.

- Kellard, N., Newbold, P., Rayner, A. and Ennew, C., The relative efficiency of commodity futures markets. *J. Fut. Mkts*, 1999, **19**, 413–432.
- Kinnear, K., The Brent WTI arbitrage: Linking the world's key crudes. Working Paper, 2002.
- Laws, J., Dunis, C.L. and Evans, B., Trading and filtering futures spread portfolios: Further applications of threshold and correlation filters. Working Paper, 2008.
- Lin, Y., McRae, M. and Gulati, C., Loss protection in pairs trading through minimum profit bounds: A cointegration approach. *J. Appl. Math. Decis. Sci.*, 2006, **6**, 1–14.
- Marshall, B.R., Cahan, R.H. and Cahan, J.M., Can commodity futures be profitably traded with quantitative market timing strategies? Working Paper, 2008.
- Mou, Y., Limits to arbitrage and commodity index investment: Front-running the Goldman roll. Working Paper, 2010.
- Papadakis, G. and Wysocki, P., Pairs trading and accounting information. Working Paper, 2007.
- Park, C.H. and Irwin, S.H., What do we know about the profitability of technical analysis? *J. Econ. Surv.*, 2007, **21**, 786–826.
- Qui, M. and Wu, Y., Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market. *J. Money, Credit, Bank.*, 2006, **38**, 2135–2158.
- Ricciardi, L.M. and Sato, S., First-passage-time density and moments of the Ornstein-Uhlenbeck process. *J. Appl. Probab.*, 1988, **25**, 43–57.
- Romano, J.P., Shaikh, A.M. and Wolf, M., Hypothesis testing in econometrics. Working Paper, 2009.
- Romano, J.P. and Wolf, M., Control of generalized error rates in multiple testing. *Ann. Statist.*, 2007, **35**, 1378–1408.
- Romano, J.P. and Wolf, M., Balanced control of generalized error rates. *Ann. Statist.*, 2010, **38**, 598–633.
- Sarr, A. and Lybek, T., Measuring liquidity in financial markets. Working Paper, 2002.
- Sato, S., Evaluation of the first-passage time probability to a square root boundary for the Wiener process. *J. Appl. Probab.*, 1977, **14**, 850–856.
- Schwarz, T.V. and Szakmary, A.C., Price discovery in petroleum markets: Arbitrage cointegration and the time interval of analysis. *J. Fut. Mkts*, 1994, **14**, 147–167.
- Shleifer, A. and Vishny, R., The limits of arbitrage. *J. Finance*, 1997, **52**, 35–55.
- Silvapulle, P. and Moosa, I., The relationship between spot and futures prices: Evidence from the crude oil market. *J. Fut. Mkts*, 1999, **19**, 175–193.
- Sullivan, R., Timmermann, A. and White, H., Data-snooping, technical trading rule performance, and the bootstrap. *J. Finance*, 1999, **54**, 1647–1691.
- Thomas, M.U., Some mean first-passage time approximations for the Ornstein-Uhlenbeck process. *J. Appl. Probab.*, 1975, **12**, 600–604.
- Trapletti, A., Geyer, A. and Leisch, F., Forecasting exchange rates using cointegration models and intraday data. *J. Forecast.*, 2002, **21**, 151–166.
- Vasicek, O.A., An equilibrium characterization of the term structure. *J. Forecast.*, 1977, **5**, 177–188.
- Vidyamurthy, G., *Pairs Trading: Quantitative Methods and Analysis*, 2004 (Wiley: Hoboken, NJ).
- Whistler, M., *Trading Pairs: Capturing Profits and Hedging Risk with Statistical Arbitrage Strategies*, 2004 (Wiley: Hoboken, NJ).
- White, H., A reality check for data snooping. *Econometrica*, 2000, **68**, 1097–1126.

Appendix A: Liquidity Measures

TABLE A.1 Liquidity Measures: WTI

Year	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12
Average 5-day turnover rate												
2003	4.05	3.01	1.78	1.08	0.69	0.54	0.54	0.44	0.38	0.33	0.27	0.25
2004	4.15	2.85	1.64	1.02	0.73	0.60	0.60	0.45	0.39	0.34	0.31	0.29
2005	4.15	2.55	1.85	1.25	0.84	0.67	0.67	0.49	0.37	0.32	0.32	0.24
2006	4.51	2.41	1.71	1.06	0.76	0.60	0.60	0.53	0.48	0.44	0.43	0.30
2007	7.39	3.17	2.20	1.44	0.95	0.74	0.74	0.44	0.37	0.28	0.26	0.21
2008	9.09	3.59	2.48	1.80	1.33	0.94	0.94	0.50	0.41	0.32	0.31	0.34
2009	10.04	3.60	2.45	1.82	1.50	1.24	1.24	0.82	0.73	0.68	0.55	0.53
2010	11.86	3.87	2.63	2.11	1.64	1.39	1.39	1.02	0.83	0.73	0.67	0.61
Average 5-day Hui-Heubel ratio												
2003	0.018	0.017	0.026	0.040	0.058	0.072	0.072	0.092	0.116	0.120	0.143	0.191
2004	0.019	0.019	0.031	0.048	0.067	0.092	0.092	0.121	0.147	0.188	0.226	0.244
2005	0.018	0.019	0.025	0.036	0.053	0.066	0.066	0.110	0.170	0.233	0.182	0.512
2006	0.015	0.018	0.023	0.035	0.050	0.065	0.065	0.076	0.093	0.102	0.123	0.191
2007	0.011	0.014	0.019	0.028	0.041	0.055	0.055	0.095	0.127	0.162	0.240	0.267
2008	0.016	0.024	0.031	0.040	0.057	0.077	0.077	0.168	0.250	0.289	0.314	0.474
2009	0.015	0.022	0.029	0.037	0.044	0.049	0.049	0.076	0.089	0.094	0.117	0.145
2010	0.006	0.011	0.016	0.020	0.025	0.030	0.030	0.042	0.048	0.054	0.061	0.073

TABLE A.2 Liquidity Measures: Brent

Year	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12
Average 5-day turnover rate												
2003	3.01	1.86	1.69	1.11	0.77	0.53	0.40	0.41	0.41	0.56	1.06	2.05
2004	2.85	2.07	1.75	1.01	0.65	0.53	0.36	0.43	0.38	0.31	0.40	2.01
2005	4.48	2.10	2.11	1.45	0.94	0.64	0.47	0.44	0.39	0.33	0.39	0.69
2006	5.81	2.09	2.12	1.44	0.96	0.73	0.54	0.47	0.48	0.42	0.43	0.52
2007	5.12	2.35	2.41	1.63	1.07	0.81	0.63	0.49	0.37	0.37	0.36	0.33
2008	7.24	2.64	2.55	2.08	1.53	1.11	0.90	0.71	0.59	0.45	0.34	0.31
2009	5.82	2.59	2.38	1.87	1.49	1.25	1.04	0.92	0.86	0.71	0.58	0.47
2010	6.19	2.69	2.26	1.94	1.71	1.50	1.29	1.11	1.02	0.89	0.73	0.67
Average 5-day Hui-Heubel ratio												
2003	0.185	0.093	0.030	0.041	0.062	0.089	0.123	0.163	0.125	0.172	0.257	0.156
2004	0.028	0.026	0.029	0.049	0.078	0.108	0.194	0.355	0.321	0.527	0.415	0.220
2005	0.013	0.022	0.021	0.029	0.048	0.078	0.129	0.262	0.340	0.903	0.355	0.651
2006	0.010	0.020	0.019	0.027	0.040	0.059	0.099	0.155	0.199	0.177	0.629	1.621
2007	0.009	0.017	0.017	0.023	0.037	0.049	0.069	0.102	0.154	0.173	0.845	0.623
2008	0.013	0.036	0.034	0.052	0.061	0.082	0.089	0.124	0.165	0.275	0.363	1.926
2009	0.016	0.026	0.028	0.033	0.042	0.048	0.058	0.069	0.073	0.092	0.108	0.155
2010	0.009	0.015	0.017	0.020	0.023	0.027	0.030	0.037	0.041	0.047	0.061	0.067

TABLE A.3 Liquidity Measures: Heating Oil

Year	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12
Average 5-day turnover rate												
2003	4.13	2.29	1.24	0.77	0.60	0.48	0.39	0.39	0.35	0.33	0.29	0.36
2004	3.84	1.98	1.12	0.73	0.60	0.46	0.38	0.35	0.35	0.37	0.29	0.30
2005	3.88	1.89	1.16	0.78	0.57	0.49	0.44	0.39	0.36	0.36	0.32	0.38
2006	4.05	1.83	1.19	0.79	0.56	0.48	0.53	0.40	0.38	0.39	0.34	0.31
2007	4.27	2.02	1.43	1.02	0.79	0.65	0.58	0.48	0.45	0.40	0.45	0.45
2008	5.09	2.39	1.73	1.32	1.06	0.87	0.82	0.75	0.66	0.49	0.41	0.41
2009	4.90	2.25	1.33	1.01	0.85	0.75	0.73	0.67	0.64	0.58	0.56	0.50
2010	6.87	2.33	1.53	1.18	1.01	0.77	0.68	0.64	0.57	0.46	0.39	0.43
Average 5-day Hui-Heubel ratio												
2003	0.020	0.026	0.041	0.064	0.084	0.110	0.122	0.134	0.184	0.317	0.517	0.466
2004	0.021	0.030	0.050	0.072	0.090	0.118	0.139	0.159	0.186	0.227	0.253	0.377
2005	0.021	0.030	0.045	0.066	0.088	0.122	0.166	0.156	0.203	0.191	0.234	0.378
2006	0.017	0.025	0.037	0.053	0.073	0.087	0.079	0.116	0.135	0.229	0.388	0.332
2007	0.013	0.020	0.027	0.036	0.047	0.059	0.064	0.094	0.089	0.121	0.143	0.228
2008	0.018	0.029	0.038	0.051	0.062	0.076	0.086	0.089	0.108	0.147	0.213	0.253
2009	0.019	0.029	0.046	0.061	0.070	0.078	0.081	0.090	0.093	0.109	0.127	0.119
2010	0.012	0.018	0.025	0.033	0.040	0.052	0.059	0.072	0.077	0.094	0.115	0.113

TABLE A.4 Liquidity Measures: Gas Oil

Year	M1	M2	M3	M4	M5	M6
Average 5-day turnover rate						
2003	5.32	1.43	1.79	0.61	0.41	0.38
2004	5.72	0.95	0.94	0.57	0.38	0.31
2005	5.49	1.54	1.32	0.72	0.45	0.38
2006	17.31	1.90	1.57	0.97	0.55	0.42
2007	11.61	2.08	1.63	0.99	0.60	0.45
2008	15.73	2.69	2.19	1.35	0.94	0.74
2009	8.29	2.54	1.76	1.11	0.78	0.64
2010	6.22	2.92	2.12	1.44	1.03	0.80
Average 5-day Hui-Heubel ratio						
2003	0.047	0.043	0.045	0.102	0.149	0.384
2004	0.149	0.169	0.090	0.116	0.171	0.330
2005	0.025	0.034	0.041	0.073	0.155	0.228
2006	0.013	0.022	0.027	0.040	0.073	0.108
2007	0.012	0.017	0.022	0.036	0.061	0.084
2008	0.017	0.033	0.037	0.064	0.088	0.115
2009	0.021	0.026	0.037	0.056	0.078	0.099
2010	0.014	0.013	0.019	0.028	0.038	0.050

Appendix B: Top 10 Profitable Trading Strategies

TABLE B.1 Top 10 Profitable Trading Strategies: 2003 and 2004

Series	Avg. daily ret. (%)	Sharpe ratio	Avg. trade length (days)	α	μ	σ	F-Test <i>p</i> -value
2003							
GO M01–HO M02 (1Y)	0.444	3.06	9.67	142.90	−0.04	0.24	0.00
WTI M04–GO M04 (1Y)	0.404	4.12	9.21	105.55	−0.09	0.24	0.00
GO M01–HO M01 (1Y)	0.373	2.31	10.44	106.44	−0.03	0.26	0.00
GO M01–HO M02 (3Y)	0.362	2.46	17.40	51.36	−0.03	0.31	0.00
Brent M11–WTI M01 (1Y)	0.355	3.32	18.64	3.76	−0.20	0.20	0.02
Brent M01–GO M03 (3Y)	0.327	3.04	18.50	17.00	−0.13	0.31	0.00
Brent M01–GO M03 (2Y)	0.327	3.03	16.19	23.87	−0.13	0.28	0.00
WTI M05–GO M04 (1Y)	0.320	3.22	10.88	76.15	−0.10	0.24	0.00
WTI M04–GO M03 (2Y)	0.319	3.01	14.50	30.24	−0.10	0.28	0.00
GO M01–HO M02 (2Y)	0.306	2.07	17.40	94.04	−0.03	0.28	0.00
2004							
GO M04–HO M04 (3Y)	0.392	3.22	12.48	89.72	−0.04	0.26	0.00
GO M05–HO M05 (3Y)	0.328	2.76	12.38	78.74	−0.04	0.26	0.00
GO M04–HO M05 (3Y)	0.310	2.57	12.48	64.79	−0.03	0.27	0.00
GO M05–HO M05 (2Y)	0.292	2.44	13.68	73.71	−0.04	0.25	0.00
GO M04–HO M04 (2Y)	0.283	2.27	17.47	82.53	−0.04	0.26	0.00
Brent M12–GO M01 (1Y)	0.281	2.33	26.10	10.89	−0.28	0.38	0.02
WTI M04–GO M06 (3Y)	0.281	2.37	17.47	28.99	−0.08	0.26	0.00
Brent M11–GO M01 (1Y)	0.274	2.28	26.10	11.28	−0.27	0.37	0.02
GO M05–HO M04 (3Y)	0.273	2.20	20.15	53.24	−0.04	0.27	0.00
WTI M05–GO M06 (3Y)	0.270	2.31	17.47	41.83	−0.10	0.25	0.00

TABLE B.2 Top 10 Profitable Trading Strategies: 2005 and 2006

Series	Avg. daily ret. (%)	Sharpe ratio	Avg. trade length (days)	α	μ	σ	F-Test p-value
2005							
Brent M12-GO M06 (3Y)	0.288	2.99	13.68	33.48	-0.21	0.24	0.00
Brent M01-GO M01 (1Y)	0.280	2.19	17.13	20.35	-0.20	0.30	0.00
Brent M11-GO M06 (3Y)	0.267	2.76	15.29	36.58	-0.21	0.24	0.00
WTI M08-GO M02 (1Y)	0.265	2.23	19.92	22.79	-0.17	0.30	0.00
Brent M03-GO M01 (2Y)	0.255	2.04	19.69	19.45	-0.20	0.31	0.00
Brent M10-GO M06 (2Y)	0.249	2.57	15.29	41.76	-0.20	0.23	0.00
WTI M06-GO M02 (3Y)	0.246	2.10	21.25	17.17	-0.14	0.29	0.00
WTI M04-HO M01 (1Y)	0.244	2.66	28.89	6.68	-0.15	0.18	0.05
GO M01-HO M01 (3Y)	0.243	1.77	19.92	42.34	-0.03	0.34	0.00
WTI M08-GO M02 (3Y)	0.242	2.03	23.55	14.30	-0.16	0.29	0.00
2006							
GO M02-HO M03 (2Y)	0.300	3.07	14.44	52.70	-0.02	0.29	0.00
GO M01-HO M02 (2Y)	0.265	2.46	18.57	39.03	-0.02	0.33	0.00
GO M02-HO M03 (3Y)	0.263	2.67	16.25	45.65	-0.03	0.30	0.00
GO M04-HO M04 (1Y)	0.238	2.51	18.57	43.18	-0.01	0.26	0.00
GO M03-HO M03 (1Y)	0.235	2.41	18.57	35.95	-0.01	0.27	0.00
WTI M02-GO M05 (1Y)	0.228	2.32	20.00	16.83	-0.19	0.26	0.00
WTI M03-GO M01 (2Y)	0.228	2.13	23.64	22.05	-0.14	0.32	0.00
WTI M02-GO M01 (2Y)	0.222	2.04	23.64	18.38	-0.14	0.33	0.00
GO M05-HO M05 (1Y)	0.220	2.37	16.25	57.80	-0.01	0.25	0.00
WTI M01-GO M05 (1Y)	0.217	2.09	23.64	18.25	-0.20	0.28	0.00

TABLE B.3 Top 10 Profitable Trading Strategies: 2007 and 2008

Series	Avg. daily ret. (%)	Sharpe ratio	Avg. trade length (days)	α	μ	σ	F-Test p-value
2007							
WTI M02–GO M01 (2Y)	0.319	2.78	16.31	34.48	−0.16	0.29	0.00
Brent M02–GO M03 (1Y)	0.309	3.16	9.00	101.02	−0.17	0.21	0.00
Brent M05–GO M03 (1Y)	0.295	3.23	10.88	94.83	−0.15	0.21	0.00
Brent M03–GO M03 (1Y)	0.295	3.10	9.00	108.19	−0.16	0.20	0.00
Brent M05–GO M04 (1Y)	0.292	3.22	11.35	96.84	−0.16	0.20	0.00
Brent M04–GO M04 (1Y)	0.288	3.11	11.35	106.38	−0.17	0.20	0.00
GO M01–HO M01 (2Y)	0.287	2.66	15.35	56.23	0.00	0.29	0.00
GO M02–HO M02 (1Y)	0.283	2.88	12.43	122.85	0.00	0.22	0.00
Brent M02–GO M01 (1Y)	0.282	2.62	11.86	79.97	−0.15	0.23	0.00
WTI M03–GO M01 (2Y)	0.278	2.50	18.64	32.09	−0.15	0.28	0.00
2008							
GO M02–HO M01 (2Y)	0.283	2.26	15.35	91.53	0.01	0.23	0.00
GO M03–HO M02 (2Y)	0.261	2.07	15.35	113.21	0.00	0.22	0.00
GO M03–HO M04 (1Y)	0.248	1.99	16.38	115.36	−0.01	0.21	0.00
GO M02–HO M01 (3Y)	0.247	1.96	20.08	67.23	0.01	0.26	0.00
GO M02–HO M01 (1Y)	0.239	1.89	17.40	176.82	0.00	0.20	0.00
GO M01–HO M01 (2Y)	0.229	1.49	18.29	116.17	0.00	0.24	0.00
GO M01–HO M01 (3Y)	0.223	1.45	21.42	70.32	0.00	0.28	0.00
GO M03–HO M05 (1Y)	0.219	1.75	21.83	48.26	−0.01	0.23	0.00

TABLE B.4 Top 10 Profitable Trading Strategies: 2009 and 2010

Series	Avg. daily ret. (%)	Sharpe ratio	Avg. trade length (days)	α	μ	σ	F-Test p-value
2009							
GO M01–HO M02 (2Y)	0.66	4.77	7.46	90.81	0.01	0.26	0.00
GO M02–HO M03 (2Y)	0.59	4.32	9.00	91.26	0.01	0.25	0.00
GO M01–HO M02 (3Y)	0.56	3.96	9.00	93.50	0.01	0.25	0.00
GO M02–HO M02 (1Y)	0.56	3.99	9.67	102.62	0.02	0.29	0.00
GO M01–HO M01 (3Y)	0.55	3.85	8.42	83.40	0.01	0.25	0.00
GO M01–HO M03 (3Y)	0.54	3.85	9.67	64.28	0.00	0.25	0.00
GO M03–HO M03 (1Y)	0.52	3.86	9.67	117.69	0.02	0.28	0.00
GO M02–HO M03 (3Y)	0.51	3.65	10.44	86.65	0.00	0.24	0.00
GO M01–HO M02 (1Y)	0.49	3.44	9.00	121.83	0.02	0.28	0.00
GO M04–HO M05 (3Y)	0.49	3.79	9.00	77.89	0.00	0.24	0.00
2010							
WTI M01–GO M05 (1Y)	0.114	1.43	65.25	14.94	−0.16	0.43	0.00
GO M05–HO M06 (1Y)	0.104	1.58	23.64	164.86	−0.01	0.27	0.00
GO M04–HO M05 (1Y)	0.103	1.55	23.64	160.37	−0.01	0.28	0.00
Brent M09–WTI M01 (3Y)	0.102	2.31	46.60	4.12	0.06	0.27	0.00
Brent M10–WTI M01 (3Y)	0.101	2.25	45.40	3.83	0.06	0.28	0.00
Brent M08–WTI M01 (3Y)	0.099	2.29	46.60	4.49	0.05	0.27	0.00
WTI M02–HO M10 (1Y)	0.083	2.31	52.20	7.75	−0.18	0.21	0.00
GO M05–HO M07 (1Y)	0.081	1.22	31.13	184.38	−0.02	0.26	0.00
Brent M05–WTI M02 (2Y)	0.070	2.13	52.20	7.23	0.03	0.17	0.01
Brent M06–WTI M02 (2Y)	0.067	1.99	52.20	5.99	0.04	0.17	0.00

II

Other Commodities

- 8 **Inversion of Option Prices for Implied Risk-Neutral Probability Density Functions: General Theory and Its Applications to the Natural Gas Market** *Yijun Du, Chen Wang, and Yibing Du* 151
Introduction • Formulation and Solution Techniques of the Option Inverse Problem • Comparison of Methods and Determination of the Implied Risk-Neutral PDF: An Example from the Natural Gas Market • Empirical Results from the Inversion of Option Prices for the September 2007 NYMEX Natural Gas Futures • Conclusions
- 9 **Investing in the Wine Market: A Country-Level Threshold Cointegration Approach** *Lucia Baldi, Massimo Peri, and Daniela Vandone*..... 173
Introduction • Wine: The Theoretical Framework • The Mediobanca Index: Evidence on Stock Performance • Econometric Methodology • Empirical Results • Conclusion
- 10 **Cross-Market Soybean Futures Price Discovery: Does the Dalian Commodity Exchange Affect the Chicago Board of Trade?** *Liyan Han, Rong Liang, and Ke Tang.*.. 191
Introduction • The Data • The SVAR Model • The VEC Model • Conclusion
- 11 **The Structure of Gold and Silver Spread Returns** *Jonathan A. Batten, Cetin Ciner, Brian M. Lucey, and Peter G. Szilagyi* 213
Introduction • Data • Rescaled Adjusted Range Analysis • Trading Implications • Conclusions
- 12 **Gold and the U.S. Dollar: Tales from the Turmoil** *Paolo Zagaglia and Massimiliano Marzo* 227
Introduction • The Dataset • Some Structural Evidence on Volatility Spillovers • A Look at the Tails • Conclusion
- 13 **A Flexible Model of Term-Structure Dynamics of Commodity Prices: A Comparative Analysis with a Two-Factor Gaussian Model** *Hiroaki Suenaga*..... 243
Introduction • Comparison of Two Modelling Approaches for the Term Structure of Commodity Futures • Data and Estimation • Implications for the Optimal Hedging Strategy • Conclusion
- 14 **Application of a TGARCH-Wavelet Neural Network to Arbitrage Trading in the Metal Futures Market in China** *Lei Cui, Ke Huang, and H.J. Cai*..... 279
Introduction • Model Development: TGARCH-WNN Statistical Arbitrage • Sample Selection • Empirical Study • Conclusion

Inversion of Option Prices for Implied Risk-Neutral Probability Density Functions: General Theory and Its Applications to the Natural Gas Market

8.1	Introduction	152
	Short Survey of the Literature on the Estimation of PDFs from Option Prices • Short Survey of the Literature on the Application of Inverse Theory in the Calibration of Option Pricing Models • Overview of the Paper	
8.2	Formulation and Solution Techniques of the Option Inverse Problem	156
	Representation of Option Prices and Parameterization of the PDF • Ill-Posedness and Constraints • Penalty Function Technique • Positivity Constraints	
8.3	Comparison of Methods and Determination of the Implied Risk-Neutral PDF: An Example from the Natural Gas Market... 160	
	Data • Implied Risk-Neutral PDFs • Resolving Power • Similarities and Differences of Solutions	
8.4	Empirical Results from the Inversion of Option Prices for the September 2007 NYMEX Natural Gas Futures	166
8.5	Conclusions.....	167
	Acknowledgements.....	168
	References.....	168

Yijun Du

Chen Wang

Yibing Du

This paper applies inverse theory to the estimation of the implied risk-neutral probability density function (PDF) from option prices. A general framework of inverting option prices for the implied risk-neutral PDF is formulated from the option pricing formula of Cox and Ross (1976). To overcome the non-uniqueness and instability inherent in the option inverse problem, the smoothness requirement for the shape of the PDF and a prior model are introduced by a penalty function. Positivity constraints are included as a hard bond on the PDF values. The option inverse problem then becomes a non-negative least-squares problem that can be solved by classic methods such as the non-negative least-squares program of Lawson and Hanson (1974). The best solution is not the one that gives the

best fit to the observed option prices or provides the smoothest PDF, but the one that gives the optimal trade-off between the goodness-of-fit and smoothness of the estimated risk-neutral PDF. The proposed inversion technique is compared with the models of Black–Scholes (BS), a mixture of two lognormals (MLN), Jarrow and Rudd’s Edgeworth expansion (JR), and jump diffusion (JD) for the estimation of the PDF from the option prices associated with the September 2007 NYMEX natural gas futures. It is found that the inversion technique not only gives the best goodness-of-fit, but also a significantly better model resolution. An empirical study for the last three months of the September 2007 futures contract shows that the shapes of the estimated PDFs become more symmetric as the futures contract becomes closer to the expiration date. The dispersion of the estimated PDFs decreases with decreasing time to expiry, indicating the resolution of uncertainty with passing time.

Keywords: Risk neutral distributions; Option pricing; Parameter estimation techniques; Goodness of fit tests; Inverse theory

8.1 Introduction

Estimating the risk-neutral probability density function (PDF) of asset prices from option prices has been an important area of research and application in finance in the last 20 years. The estimated PDF can shed light on information concerning market expectations of future behaviour of underlying asset prices, and therefore is highly useful in risk management, derivatives pricing and monetary policy making.

The price of a European call option with strike K at expiry T can be written in terms of the risk-neutral PDF, $q(p)$, for the underlying price as

$$c(K) = e^{-rT} \int_0^{\infty} \max[0, p - K] q(p) dp \quad (8.1)$$

(Cox and Ross 1976). By differentiating Equation (8.1) twice, $q(p)$ can be obtained as

$$\frac{\partial^2 c(K)}{\partial K^2} = e^{-rT} q(p) \quad (8.2)$$

(Breedon and Litzenberger 1978). Note that the derivation of both Equations (8.1) and (8.2) makes no assumptions about the underlying asset price dynamics, as long as markets are efficient. In Equation (8.2), $c(K)$ assumes to be twice differentiable. However, this assumption is not necessary for estimating a discretized PDF using Equation (8.2). The classic Black–Scholes (1973) option pricing model assumes $q(p)$ to be lognormal and therefore, the underlying asset price follows a stochastic process, the so-called geometric Brownian motion (GBM). Even though most of the time asset prices follow GBM, they display small but important ‘non-log normality’. In particular, the frequency and direction of large moves in asset prices can be quite different than predicted by the GBM model. The implied volatilities calculated from the observed option prices based on the Black–Scholes model vary with strike (a phenomenon called volatility smile or skew), indicating that market participants make more complex assumptions than GBM about the distribution of the underlying asset price. This paper applies inverse theory to the estimation of the implied risk-neutral PDF from the observed option prices. In the remainder of the introduction, we first review existing methods for estimating implied risk-neutral PDFs, then review the previous applications of inverse theory in calibrating option pricing models, and finally provide an overview of the paper.

8.1.1 Short Survey of the Literature on the Estimation of PDFs from Option Prices

Various methods have been suggested to estimate non-GBM risk-neutral PDFs from option prices with examples of their applications to particular markets. Cont (1997) and Bahra (1997) provided two earlier

surveys of various methods and applications. Cont (1997) focused on particular methods for recovering risk-neutral PDFs, whereas Bahra (1997) reviewed both methods and applications. Cont (1997) classified the various methods into expansion, non-parametric and parametric methods. We use related, but different, categories and classify the various methods and applications into the following three broad categories.

1. By assuming a particular stochastic process to account for the deviation from log normality for the process of the underlying asset price, the observed option prices are used to estimate the parameters of the assumed stochastic process with the implied risk-neutral PDF obtained from the estimation process (e.g. Malz 1995, Bates 1996). Further modification to the non-GBM stochastic process is to model the stochastic volatility (e.g. Heston 1993, Cont and Tankov 2004a). The estimation of risk-neutral PDFs then involves time-consuming numerical integral (Jondeau and Rockinger 2000).
2. Several methods have been developed based on Equation (8.1) to account for the non-GBM PDFs, typically using a non-linear optimization method to find the exact forms of the PDF that minimizes the sum of squared errors between the calculated and observed option prices. These methods differ in the amount of structure they place on the PDF to be estimated and can be classified into three groups. The first group of methods assumes a particular functional form for the PDF, for example a mixture of lognormal densities (e.g. Ritchey 1990, Bahra 1997, Melick and Thomas 1997) or a non-lognormal PDF such as the Burr III distribution (Sherrick *et al.* 1996). The second group of methods is lognormal-based expansion techniques in which a lognormal PDF is assumed as a zero-order approximation of the true PDF, then non-log normality is accounted for by higher-order expansions such as the Hermite polynomial expansion (Madan and Milne 1994, Abken *et al.* 1996) and the Edgeworth expansion (Jarrow and Rudd 1982, Corrado and Su 1996). In most practical applications the expansions have been limited to fourth order. The third- and fourth-order terms have a direct correlation with the skewness and Kurtosis, respectively. The biggest disadvantage of the expansion techniques is that the polynomial expansions can yield negative PDF values for certain combinations of the skewness–kurtosis parameters. Therefore, the positivity constraints require that the expansion approach be used for relatively moderate departures from log normality (e.g. Rubinstein 1998, Jondeau and Rockinger 2001). The third group of methods is maximum entropy approaches (Buchen and Kelly 1996, Stutzer 1996). The idea is to choose, between all PDFs that satisfy Equation (8.1), the one with maximum entropy. The estimated PDF is the one closest to the prior PDF q_0 (q_0 can be a pre-selected PDF or a historical density). However, the estimated PDF is highly sensitive to the chosen prior. Most importantly, the estimation problem remains ill-posed since entropy is only used as a model selection criterion without regularization (e.g. Cont and Tankov 2006).
3. Several techniques for estimating the risk-neutral PDF based on Equation (8.2) have also been developed. By interpolating the observed option prices and fitting them to some functional form, $\varphi(K)$, the risk-neutral PDF can be obtained by differentiating $\varphi(K)$ twice according to Equation (8.2) (e.g. Shimko 1993, Longstaff 1995, Malz 1997). Jackwerth and Rubinstein (1996) suggested a maximum smoothness criterion that essentially uses a butterfly spread variant of Equation (8.2) to minimize the curvature of the estimated implied PDF. Ait-Sahalia and Lo (1998) applied the Nardaraya–Watson non-parametric kernel regression to estimate the functional form that relates the call option prices to the strike price, and then obtained the risk-neutral PDF according to Equation (8.2).

Several studies have been performed to compare some of the methods for estimating the risk-neutral PDF (e.g. Bahra 1997, Campa *et al.* 1998, Jackwerth 1999, Jondeau and Rockinger 2000). It is found that the values for the first two moments of the PDF are similar across various models and that there exist large differences in the higher moments of the PDF. Although each method of estimating the risk-neutral

PDF has its advantages and drawbacks, the main difference between different methods is the extent to which they constrain the shape of the PDF to be estimated. At one extreme, Longstaff (1995) imposes no constraints, but the result can be a rather rough or spiky distribution. At the other extreme, Rubinstein (1994) and Jackwerth and Rubinstein (1996) constrain their distributions to be those with the smallest possible deviation from the lognormal with the maximum smoothness. Therefore, measures of skewness and kurtosis are highly model-dependent.

The abundance of methods for estimating the risk-neutral PDF is indeed due to the inherent non-uniqueness and instability associated with the inversion of option prices for the PDF: there are infinite solutions for estimating the PDF from the limited option prices observed. How do we determine a solution that is adequate and meaningful for the specific estimation problem? What are the proper constraints on the shape of the PDF to be estimated? How much smoothness is needed for the PDF? In this paper, we will answer these questions by applying inverse theory.

8.1.2 Short Survey of the Literature on the Application of Inverse Theory in the Calibration of Option Pricing Models

Inversion techniques have been widely applied to solve various practical problems in fields such as geophysics (e.g. Backus and Gilbert 1967, Wiggins 1972, Tarantola and Valette 1982), ocean and atmospheric sciences (Bennett 2002) and medical imaging (Natterer 2001). Since inverse problems are often formulated in infinite-dimensional space, limitations to a finite number of observational data will typically lead to two major problems in inverse problems: non-uniqueness and instability. *A priori* information or constraints are required to eliminate the non-uniqueness and instability of the solution (e.g. Backus and Gilbert 1970, Franklin 1970, Jackson 1979). Different ways of incorporating constraining information result in different approaches to the investigation and numerical resolution of inverse problems (for overviews, see Parker 1977, Tikhonov and Arsenin 1977, Menke 1984, Hofmann 1986, Baumeister 1987, Tarantola 1987, Engl *et al.* 1996, Hansen 1998).

The applications of inverse theory in finance are relatively more recent and mainly in calibrating volatility surfaces (both volatility smile and term structure) from option prices in the finance literature (for overviews, see Bouchouev and Isakov 1999, Cont and Tankov 2004a, Engl 2007, Kindermann and Pikkarainen 2009). These applications can be classified into the following three categories.

1. In the first category of applications, the stochastic diffusion equations for the underlying assets are reduced to simpler and solvable parabolic-type partial differential equations under some strong assumptions (Bouchouev and Isakov 1997, 1999, Berestycki *et al.* 2000, 2002, 2004). By assuming options are close to expiration or deep in- or out-of-the money, Berestycki *et al.* showed that there is a partial differential equation of degenerate parabolic type linking the local and implied volatilities and that asymptotic expressions for the implied volatility can be derived. By assuming a pure spot-price dependence for volatility ($\sigma(s, t)$ is independent of t), Bouchouev and Isakov reduced the stochastic diffusion equation to an inverse parabolic problem and obtained a non-linear Fredholm integral equation that can be solved iteratively for unknown spot-price-dependent $\sigma(s)$ after dropping terms of the higher-order time to maturity. Bouchouev and Isakov's approach is applicable for finding short-term volatility, but not suitable for estimating long-term volatility.
2. In the second category of applications, the smoothness of the function is used to regularize an ill-posed inverse problem (Jackwerth and Rubinstein 1996, Lagnado and Osher 1997, Coleman *et al.* 1999). Jackwerth and Rubinstein's maximum smoothness criterion finds a solution that minimizes the curvature of the PDF, whereas Lagnado and Osher find a solution that minimizes the gradient of the local volatility function. By assuming that the local volatility function is smooth and recognizing that splines can approximate smooth curves and surfaces, Coleman *et al.* (1999) suggested using a spline with a fixed set of spline knots and end condition to approximately

represent the local volatility function. The volatility values at knots are then determined by solving a small nonlinear optimization problem. However, they had to assume the number of spline knots to be equal to or less than the number of observations in order for the estimation problem to be well-posed. In practical applications, the number of observations is often small. Therefore, the number of spline knots would normally exceed the number of observations in order for a spline to represent the true volatility function well.

3. In the third category of applications, a prior model has been used to regularize the ill-posed volatility calibration problems. Depending on how a prior model is introduced, these applications can be classified into two groups. (i) The first group uses relative entropy to introduce a prior model into regularization (e.g. Avellaneda *et al.* 1997, Cont and Tankov 2004a,b, 2006). Avellaneda *et al.* assumed a diffusion process and a prior diffusion, and then solved for a solution that has the smallest entropy distance to the prior diffusion using a dynamic programming approach. However, as Cont and Tankov (2004b) pointed out, the problem of minimizing the relative entropy distance to a prior without regularization is still not well-posed. In calibrating jump and exponential Lévy models, Cont and Tankov (2004a,b, 2006) regularize the calibration problem by reformulating it as the problem of finding a risk-neutral Lévy measure L that minimizes $J(L) = \|C(L) - C_\delta\|^2 + \lambda E(L/L_0)$, which is the sum of the pricing error and the relative entropy with respect to a prior Lévy measure L_0 . Here $C(L)$ is the option price calculated by Lévy measure L , C_δ is the possibly noisy observed option price, E defines the relative entropy of L with respect to L_0 , and λ is the regularization parameter. The relative entropy is thus used as a penalization term to regularize the ill-posed calibration problem. (ii) If the relative entropy term $E(L/L_0)$ is changed to be the L_2 norm of $(L - L_0)$ as $\|L - L_0\|^2$, then it leads to the second group of methods for introducing a prior model, i.e. the classical Tikhonov regularization (Crépey 2003a,b, Kindermann and Pikkarainen 2009). In both relative entropy and Tikhonov regularizations, the regularization parameter λ plays a trade-off between the accuracy of calibration and model uncertainty. A very large λ will result in low accuracy of calibration and bias the solution towards the prior model. On the other hand, a very small λ will result in high accuracy of calibration but cause large model uncertainty due to insufficient regularization. λ has been determined using the so-called Morozov discrepancy principle (e.g. Cont and Tankov 2004a, Kindermann and Pikkarainen 2009). The Morozov discrepancy principle (Morozov 1984) will guarantee a regularized solution. However, the solution may not be optimal and is highly dependent on the choice of the acceptable maximum error. Another problem with the relative entropy or Tikhonov approach is that it finds a solution that is closest to the chosen prior model. Therefore, the choice of a prior becomes very crucial since an improper prior model may bias the solution, especially when data are scarce or missing. This is true for any method that uses a prior model to regularize an ill-posed inverse problem.

8.1.3 Overview of the Paper

Section 8.2 provides a general framework for the formulation and solution techniques of option inverse problems for estimating the implied risk-neutral PDF. We do not make any assumptions about the underlying stochastic process nor about the particular form of the PDF. Instead, the Cox and Ross Equation (8.1) is discretized directly into a set of ill-posed linear equations with limited option prices. To regularize the ill-posed option inverse problem, smoothness constraints on the shape of the PDF to be estimated and a prior model for the PDF are introduced by a penalty function. By adding positivity constraints to ensure meaningful PDF values, we reformulate the ill-posed option inverse problem as a non-negative least-square problem that can be solved very easily. In contrast to the methods of determining the regularization parameter in the existing financial literature (which may not yield optimal regularization), we apply an L-curve to determine the optimal regularization parameter. We also examine the resolving power of the proposed regularization method by calculating the model resolution matrix. Section 8.3 shows an example of applying the proposed inverse techniques to estimate the implied risk-neutral PDF from

the options observed on August 28, 2007 for the September 2007 NYMEX natural gas futures. For this example, we only use smoothness for regularization in order to avoid the possible bias caused by a prior model. The estimated PDFs are compared with those derived from the popular models of Black–Scholes, a mixture of two lognormals, the Edgeworth expansion and jump diffusion. The models are also compared based on their power to resolve the true PDF. It is found that our inversion technique not only gives the best goodness-of-fit to the observed option prices, but also has significantly better resolving power. [Section 8.4](#) shows the empirical results by applying the inversion model to the estimation of implied PDFs from the daily option closing prices for the last 3 months of the September 2007 natural gas futures. It shows that the estimated PDFs become more symmetric and the dispersion of the estimated PDFs decreases as the futures contract becomes closer to the expiration date. [Section 8.5](#) concludes.

The contributions of this paper are threefold. First, we directly discretize the Cox and Ross Equation (8.1) into a linear ill-posed operator and no assumptions about the underlying stochastic process or the specific form of a PDF are made. This is in contrast to most studies, which assume a particular stochastic process or form of the estimated PDF. Second, smooth to a certain degree for the shape of the estimated PDF is the only assumption made for the regularization of the ill-posed option inverse problem. No prior model is assumed in our example. The regularization parameter is determined optimally by an L-curve. Thus, the best solution is the one that gives the optimal trade-off between the goodness-of-fit and the smoothness of the estimated PDF, not the one that gives the smoothest PDF as formulated by previous studies. Finally, a matrix called a model resolution matrix is calculated to examine the resolving power of the proposed inverse technique and several other methods.

8.2 Formulation and Solution Techniques of the Option Inverse Problem

8.2.1 Representation of Option Prices and Parameterization of the PDF

Our forward formulation is the Cox and Ross integral of Equation (8.1). We parameterize the continuous probability density function (PDF) along the p -axis with a series of linear segments ([Figure 8.1](#)).

The PDF within each element is assumed to be uniform. The integral (8.1) can be discretized into a set of linear equations. The discrete approximation of the observed option prices then can be written as

$$c(K_i) = e^{-rT} \sum_{j=1}^m q_j(p_j) (p_{j+1} - p_j) \max(p_j - K_i, 0), \quad (8.3a)$$

$$p(K_i) = e^{-rT} \sum_{j=1}^m q_j(p_j) (p_{j+1} - p_j) \max(K_i - p_j, 0), \quad (8.3b)$$

with the following equality constraint on $q_j(p_j)$:

$$1 = \sum_{j=1}^m q_j(p_j) (p_{j+1} - p_j), \quad (8.3c)$$

where $c(K_i)$ and $p(K_i)$ are the observed option prices for the call and put options with the strike of K_i , respectively, and $q_j(p_j)$ is the PDF value at the asset price of p_j . Repeated index j implies summation from the element 1 to m . Equation (8.3) can be written in matrix form as

$$C = GQ, \quad (8.4)$$

where C is the vector of n option prices plus the numerical value of 1, G is an $(n+1) \times m$ matrix of discrete coefficient functions, and Q is the vector of m unknown densities. As long as the discretization is relatively fine enough, the discrete approximation can then adequately represent the true, continuous, probability density function, implying that the discrete linear elements must be reasonably small. In




FIGURE 8.1 Discretization of a probability density function. The PDF is discretized into m linear elements of width $\Delta_j = p_{j+1} - p_j$ for the j th element. The PDF within each element is assumed to be uniform.

general, this means that the number of model parameters, m , exceeds the number of option prices, n [$m > n$]. Therefore, the system of equations (8.4) is underdetermined and the solution is not unique.

8.2.2 Ill-posedness and Constraints

The first principle of inversion theory is the notion of a well-posed problem and is commonly attributed to Hadamard (1923). A problem characterized by the operator equation $Ax = y$ is considered to be well-posed if the following three conditions are satisfied: (1) the solution x exists for any y (existence); (2) there exists only one solution (uniqueness); and (3) small perturbations in y produce small perturbations on the solution x without imposing additional constraints (stability). Otherwise, the problem is said to be ill-posed (or improperly posed). For example, an overdetermined problem (number of data > number of unknowns) violates condition (1), an underdetermined problem violates condition (2) and possibly (3), and an even-determined problem may violate condition (3). It is evident from the previous section that the system of equations (8.4) is ill-posed. Clearly, no analysis is conclusive unless a well-posed character is restored between physically meaningful quantities. Given an ill-posed problem, *a priori* information or constraints should be added to the system of equations (8.4) in order to find a particular solution (Lanczos 1961, Backus 1970, Franklin 1970, Jackson 1972, 1979).

Constraints play their most important role in restoring the well-posed character of the existence, uniqueness, and stability for a physically meaningful solution. There are two classes of constraints. (1) Objective constraints are any law of nature or any technical confinement that prescribes constraints in order for the elements to be identified or for the function to be found. This kind of constraint has the highest degree of reliability. For example, PDF values must be positive. (2) Subjective constraints include all other kinds of constraints. For example, many functions arising in nature are smooth to a certain

degree. Therefore, one can expect the solution that we search for to be smooth to some degree, at least over subintervals. There are also two forms of constraints (Backus 1988): a ‘soft’ bound is a probability distribution P_x on the model space X that describes the observer’s opinion about where the model x is likely to be in X ; and a ‘hard’ bound on x is an inequality.

8.2.3 Penalty Function Technique

Following Du *et al.* (1992) we introduce constraints by the penalty function technique. It is assumed that the constraints on the model parameters are generally given by $HQ = d_0$, and a prior model Q_0 . The problem to be solved is to find the density vector, Q , that minimizes the penalty function

$$F(Q) = \|GQ - C\|^2 + \alpha^2 \|Q - Q_0\|^2 + \beta^2 \|HQ - d_0\|^2, \quad (8.5)$$

where the first term is the norm of the residuals between calculated and observed option prices, α^2 is a damping factor that weights the prior model, and β^2 is a penalty factor that weights the constraints on the model. Both α^2 and β^2 are regularization parameters. If $\beta^2 = 0$, Equation (8.5) reduces to classic Tikhonov regularization (e.g. Tikhonov and Arsenin 1977, Engle *et al.* 1996). The vector Q is a minimum of $F(Q)$ if and only if

$$(G^T G + \alpha^2 I + \beta^2 H^T H)Q = G^T C + \alpha^2 Q_0 + \beta^2 H^T d_0 \quad (8.6)$$

(Baumeister 1987). Here I is a unit matrix. The solution then is

$$Q_e = (G^T G + \alpha^2 I + \beta^2 H^T H)^{-1} (G^T C + \alpha^2 Q_0 + \beta^2 H^T d_0). \quad (8.7)$$

In this study, the solution is desired to be smooth to some degree. Particularly, we take H to be the finite difference approximation of the Laplacian operator, $\partial^2/\partial p^2$, as applied to the risk-neutral PDF, $q(p)$. Hence, we have $d_0 = 0$. If the prior model Q_0 is assumed to be zero and no constraints on H are added, then Equation (8.7) becomes the standard damped least-square solution (see, e.g., Levenberg 1944 and Marquardt 1963). The model resolution, R , is defined as the mapping between the ‘true PDF’, Q , and the estimated PDF, Q_e , i.e.

$$Q_e = RQ. \quad (8.8)$$

Then the model resolution is given by

$$R = (G^T G + \alpha^2 I + \beta^2 H^T H)^{-1} G^T G. \quad (8.9)$$

If $R = 1$, then each model parameter is uniquely determined. If R is not a unit matrix, then the estimates of the model parameters are really weighted averages of the true model parameters. The data resolution, N , that maps the observed option prices into the calculated option prices, is

$$N = G(G^T G + \alpha^2 I + \beta^2 H^T H)^{-1} G^T. \quad (8.10)$$

8.2.4 Positivity Constraints

Since PDF values must be positive, this provides the positivity constraints on the model vector Q :

$$q_i \geq 0, \quad i = 1, \dots, m. \quad (8.11)$$

However, the positivity constraints do not guarantee the uniqueness of the solution if the problem $C = CQ$ lacks a unique solution (Lawson and Hanson 1974). The first step to solve this problem, therefore, is still to regularize the ill-posed problem. One can apply the penalty function in Equation (8.5) for regulation and then the non-negative solutions can be found by solving Equation (8.6) subject to the positivity constraints. In this paper, we apply Lawson and Hanson's non-negative least-squares (NNLS) method (Lawson and Hanson 1974) to find the non-negative solutions. For easy and fast implementation of the Lawson and Hanson solution method, we can rewrite Equation (8.6) in the following form of an overdetermined system:

$$\begin{pmatrix} G \\ \alpha I \\ \beta H \end{pmatrix} Q = \begin{pmatrix} C \\ \alpha Q_0 \\ \beta d_0 \end{pmatrix}, \quad (8.12)$$

where α and β are positive (note that the normalization of Equation (8.12) becomes Equation (8.6)). Therefore, using Equation (8.12) together with the positivity constraints, we have reformulated the ill-posed inverse problem as a non-negative least-square problem that can then be solved immediately by Lawson and Hanson's NNLS program.

Since the non-negative least-square problem is non-linear, the model resolution matrix R for the solution with positivity constraints cannot be calculated by a simple matrix operation similar to Equation (8.9). A classic approach is to calculate the resolution matrix by running a forward model with a Kronecker delta density function, $\delta(p) = 1$, at the specified element to produce the synthetic option prices at the observation points, and then invert the synthetic option prices. This gives the column of R corresponding to the specified element. The matrix R is found by calculating all columns of R for all elements. When additional constraints are available, a better model resolution is guaranteed.

The regularization parameter α or β plays the role of a trade-off parameter between fitting and model uncertainty, or between fitting and roughness of the PDF (given by $\|HQ\|^2$). It has been shown that, with decreasing α or β , the total sum of squared residuals decreases, the roughness of the function to be estimated increases and the uncertainties in the estimated model parameters increase (Du *et al.* 1992). It is the main dilemma in solving ill-posed problems that the better the mathematical model describes the ill-posed problem, the worse is the 'condition number' of the associated computational problems (i.e. the more sensitive to errors). Therefore, α or β should be selected in such a way that data fitting is acceptable, and, at the same time, the roughness and uncertainty are sufficiently small. As discussed above, choices of the regulation parameter used in the existing financial literature tend to cause either over- or under-regularization and do not necessarily yield an optimal trade-off between the goodness-of-fit and smoothness of the PDF. It has been demonstrated in geophysics and other fields that the L-curve (i.e. the trade-off curve) can be easily applied to determine the optimal regularization (Du *et al.* 1992, Hansen 1992, 1998). We therefore, also select the optimal value of the trade-off-parameter using the L-curve criterion for its simplicity and convenience in this paper.

When a prior model is introduced for regulating an ill-posed inverse problem, the shape of the resultant PDF is largely determined by the shape of the prior model. Therefore, α is normally selected as zero unless there is a strong belief towards a particular prior model. The advantage with our approach is that any parametric model can be used as a prior model. Both a prior model and smoothness can be included to regulate an ill-posed inverse problem. However, too much regulation will cause the goodness-of-fit (to the observed option prices) to deteriorate. The best approach is to achieve the regulation of the ill-posed inverse problem with the least amount of constraints. In this study, we particularly examine the case where the inverse problem is solved through a penalty function with smoothness and positivity constraints (INVSP) ($\alpha = 0$ and $d_0 = 0$).

8.3 Comparison of Methods and Determination of the Implied Risk-Neutral PDF: An Example from the Natural Gas Market

8.3.1 Data

To test and compare the proposed inverse models with other popular methods, we estimated the implied risk-neutral PDFs from the option prices observed on August 28, 2007 for the September 2007 NYMEX natural gas futures. The option prices are from Bloomberg. Table 8.1 shows the closing prices of the call and put options for the strikes around the futures settlement price of \$7.89/MMBTu on June 8, 2007. The time to expire is 81 days.

The exchanged-traded natural gas options are American-style options and can be exercised early. However, early exercise is not an optimal choice for NYMEX natural gas options since they are subject to futures-type settlement (i.e. mark-to-market). As Natenberg (1994) and Oviedo (2006) pointed out, in futures markets where the options are subject to future-type settlement, there is never any economic justification for early exercise since early exercise does not result in additional cash flow. A trader will always be better off either by holding the option for more upside or selling it because of the bearish outlook, rather than exercising it early. In addition, the majority of the natural gas options are traded in the over-the-counter (OTC) market where natural gas options are European and cash-settled. These OTC options can also be settled through CME ClearPort as well as for trading on the New York trading floor. Therefore, although NYMEX natural gas options are American options, there is no early exercise premium in the option prices and holders of options do not exercise early. The NYMEX options can thus be modelled and priced as European options.

8.3.2 Implied Risk-neutral PDFs

The goodness-of-fit is measured by the root-mean-square error (RMSE), which is defined as $\left[1/n \sum_{i=1}^n (o_i - d_i)^2\right]^{1/2}$, where o_i and d_i are the i th observed and calculated option prices, respectively. The trade-off or L-curve between the RMSE and the roughness of the estimated PDF for the INVSP model is plotted in Figure 8.2, showing that the RMSE decreases with decreasing penalty factor β . This

TABLE 8.1 Closing Prices of Call and Put Options for the Nymex September 2007 Natural Gas Futures Contract on 8 June 2007. The Futures Contract is Closed at \$7.89/MMBTu. The Option Contracts Will Expire in 81 Days on 28 August 2007

Strike	Call	Put
7.60	NA	0.581
7.65	NA	0.606
7.70	NA	NA
7.75	0.795	0.657
7.80	0.772	0.683
7.85	0.749	0.709
7.90	0.727	0.737
7.95	0.706	NA
8.00	0.685	0.794
8.05	0.666	0.824
8.10	0.646	0.854
8.15	0.627	0.885
8.20	0.609	NA
8.25	0.591	NA

trade-off curve suggests that $\beta = 0.01$ is the optimal value. When β is greater than the optimal value, the fit is improved slightly, but the roughness of the estimated PDF increase greatly.

The estimated PDF for the best INVSP model is shown in Figure 8.3. Also plotted are the estimated PDFs from other models including a log-normal (BS), a mixture of two lognormals (MLN), a jump diffusion (JD) and Jarrow and Rudd's Edgeworth expansion (JR). Table 8.2 shows the RMSE for the five




FIGURE 8.2 Trade-off curve (L-curve). Root-mean-square error as a function of the roughness of the estimated PDF.




FIGURE 8.3 Implied risk-neutral PDFs estimated using different models. The PDFs are sampled at 104 discrete asset values. The option prices used for all PDFs are dated June 8, 2007, at closing for the September 2007 NYMEX natural gas futures, which settles at \$7.89/MMBTu.

models. Figure 8.4 compares the pricing errors across the strikes for the five models. The BS model has the worst fit. MLN yields a roughly 20% smaller RMSE than that given by the BS model. The JR model has the third best fit with a 50% smaller RMSE than that given by the BS model. The JD model has the second best fit with a RMSE about eight times smaller than that of the BS model. The inversion model, INVSP, has the best fit to the observed option data with a RMSE about 22 times smaller than that of the BS model. The pricing errors for the INVSP model are consistently the smallest across different strikes, while the pricing errors for all other models vary with strike. The MLN model yields a PDF with the lowest peak and fattest left tail, while the JR model has a PDF with a right tail similar to the BS model but a bumpy and negative left tail. The bumpy and negative left tail associated with the JR model is due to its incapability of capturing large skewness and kurtosis. Therefore, we will not include the JR model in the further comparison of the various models. The PDF from the INVSP model has the highest peak, a tighter distribution in the centre and a fatter right tail.

8.3.3 Resolving Power

We can see from Table 8.2 that the fit for INVSP is most satisfactory. However, other models are perhaps also acceptable since a RMSE of 0.5 cents is most of the time within the bid–ask spread for natural gas option prices. Hence, while the goodness-of-fit is the first criterion to determine whether a solution is acceptable, it is apparently not sufficient to consider only the goodness-of-fit for the selection of an estimation technique and/or for the meaningful estimation of a PDF. Since the risk-neutral PDF is not uniquely determined, it is important to identify particular features of the PDF that are well resolved by




FIGURE 8.4 Pricing errors across strikes for the five models. The option prices used are dated 8 June 2007, at closing for the September 2007 NYMEX natural gas futures, which settles at \$7.89/MMBTu.

TABLE 8.2 Root-Mean-Square Errors (\$/MMBTu) for The Five Different Models. The PDFs are Sampled at 104 Discrete Asset Values. The Option Prices Used for All PDFs are Dated June 8, 2007, At Closing for The September 2007 NYMEX Natural Gas Futures, Which Settles at \$7.89/MMBTu

INVSP	BS	JD	MLN	JR
0.000233	0.005155	0.000626	0.004149	0.002662

the observed option prices. The existing comparisons of different PDF estimation methods have focused on the goodness-of-fit, the shape of the resultant PDFs, and the robustness and ease of calculation (e.g. Campa *et al.* 1998, Jondeau and Rockinger 2000). It is shown that each method has its strengths and weaknesses. The main difference between different methods is the capability of capturing the PDF's high-order features such as skewness and kurtosis. Since the true PDF is not known, it is therefore impossible to know which method is better.

The model resolution matrix R for an estimation method provides a unique and powerful tool to evaluate the particular method's capability of recovering the true PDF from the observed option data. The diagonal terms of the resolution matrix R show the resolved amplitude of the true model parameters. The columns of resolution matrix R show how true the density at strike Q appears in the estimated density Q_e and accordingly both the resolved amplitude and the length of the estimated probability density can be determined.

To compare in detail the resolving power by different methods, we calculate the model resolution for five selected strikes at \$6.03, \$6.94, \$7.98, \$9.15, and \$10.45 per MMBTu, respectively. The calculation is as follows. A Kronecker delta density is placed at a selected strike and then both the call and put prices are calculated for the strikes from \$7.75 to \$8.25 with an increment of \$0.05. The calculated option prices are then used to estimate the implied PDFs by applying different techniques. If the estimation is unique, then a Kronecker delta density should be recovered. Otherwise, only a linear combination of the adjacent PDF values can be determined.

Table 8.3 shows the diagonal values of the resolution matrix for the selected strikes for the INVSP, BS, JD, and MLN models. The following results can be obtained from Table 8.3. (1) The resolution at the strikes with observed option prices is the best, and the resolution power deteriorates when the Kronecker delta density moves further away from the strikes where option prices are observed. (2) The resolving power is different for different techniques. INVSP has the best resolution power, far better than those by other techniques. MLN has the second best resolution power. BS and JD have the worst but similar resolution power.

Figures 8.5 through 8.9 illustrate the columns of the resolution matrix corresponding to the five selected strikes. The resolution length, determined qualitatively from the spread of the distribution, is summarized in Table 8.4. It is shown that the densities are better resolved both in length and amplitude when the strike is closer to the option observation points. INVSP has the highest resolution power and results in symmetrical PDFs with peaks matching with the locales of the Kronecker delta densities. All other methods yield asymmetrical PDFs with peaks not matching with the locales of the Kronecker delta densities. When the selected strike is inside the option observation points, INVSP's resolution length is short at \$0.40, while all other methods have a resolution length of about \$1.00 or more. The resolution length increases up to \$5.20 for PFSP, \$9.00 for MLN and to more than \$15 for BS and JD when the selected strikes are further away from the option observation points. Therefore, the PDF basically cannot be resolved both in amplitude and length when the density is

TABLE 8.3 Diagonal Values of the Resolution Matrix for Selected Strikes from Four Different Models. the Option Prices are Observed for the Strikes from \$7.75 to \$8.25 with an Increment of \$0.05

Strike	INVSP	BS	JD	MLN
\$6.03	0.0523	0.0121	0.0102	0.0379
\$6.94	0.1108	0.0221	0.0195	0.0513
\$7.98	0.8685	0.2922	0.2273	0.2055
\$9.15	0.1017	0.0129	0.0129	0.0897
\$10.45	0.0473	0.0047	0.0046	0.0211




FIGURE 8.5 Resolving kernels for a Kronecker delta density (source) located at \$6.03. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from \$7.75 to \$8.25 with an increment of \$0.05.




FIGURE 8.6 Resolving kernels for a Kronecker delta density (source) located at \$6.94. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from \$7.75 to \$8.25 with an increment of \$0.05.




FIGURE 8.7 Resolving kernels for a Kronecker delta density (source) located at \$7.98. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from \$7.75 to \$8.25 with an increment of \$0.05.




FIGURE 8.8 Resolving kernels for a Kronecker delta density (source) located at \$9.15. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from \$7.75 to \$8.25 with an increment of \$0.05.




FIGURE 8.9 Resolving kernels for a Kronecker delta density (source) located at \$10.45. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from \$7.75 to \$8.25 with an increment of \$0.05.

TABLE 8.4 Resolution Length for Selected Strikes from Four Different Models. The Option Prices are Observed for the Strikes from \$7.75 to \$8.25 with an Increment of \$0.05

Strike	INVSP	BS	JD	MLN
\$6.03	\$5.00	>\$15.00	>\$15.00	\$6.50
\$6.94	\$2.20	\$12.00	\$9.00	\$5.70
\$7.98	\$0.40	\$1.00	\$1.30	\$1.40
\$9.15	\$2.50	\$12.00	\$12.00	\$3.50
\$10.45	\$5.20	>\$15.00	>\$15.00	\$9.00

far away from the option observation points for the BS, JD and MLN models. INVSP can resolve not only the significant details inside the option observation points due to its short resolution length, but also long-wavelength features for the strikes away from the option observation points. BS, JD and MLN can only resolve the long-wavelength features even if the strikes are inside the option observation points.

8.3.4 Similarities and Differences of Solutions

It was shown in the previous section that PDFs estimated from the BS and JD models are basically similar, whereas the INVSP and MLN methods yield somewhat different PDFs, especially in the tails. These similarities and differences can be understood by the algebraic theory of generalized inverses (Nashed 1987). One generalized inverse A_{g1}^{-1} is related to any other generalized inverse A_{g2}^{-1} by the model resolutions R_1 and R_2 , and the data resolutions N_1 and N_2 (Nashed 1987):

$$A_{g1}^{-1} = (I + R_1 - R_2) A_{g2}^{-1} (I + N_2 - N_1) \quad (8.13)$$

If model resolution R_1 is similar to R_2 , and data resolution N_1 is similar to N_2 , A_{g1}^{-1} is then also similar to A_{g2}^{-1} . Hence the two solutions will be similar. For an ill-posed problem, the data are usually well resolved and consequently $N_2 - N_1 \approx 0$. Then the differences between the solutions depend primarily on the differences in model resolutions, $R_1 - R_2$. Additional constraining information improves the model resolution and accordingly resolves different and improved PDFs.

8.4 Empirical Results from the Inversion of Option Prices for the September 2007 NYMEX Natural Gas Futures

To further validate the proposed inversion technique, we applied INVSP to the estimation of the implied PDFs from the daily option closing prices for the last 3 months of the September 2007 NYMEX natural gas futures. Table 8.5 shows the average RMSEs over the 3 months for the INVSP, BS, JD and MLN models. It can be seen that the BS, JD and MLN models all yield similar average RMSEs with 0.59 cents for MLN, 0.6 cents for JD and 0.7 cents for BS. However, the average RMSE for INVSP is only 0.27 cents. Therefore, INVSP is a far better model than the other models. In the following, we will show how the shapes of the implied PDFs estimated from the INVSP model change with time.

Figure 8.10 shows the 5th, 25th, 50th, 75th and 95th percentiles of the estimated PDFs as a function of time. Also plotted are the daily NYMEX closing prices. It can be seen that the dispersion of the estimated PDFs around the 50th percentile decreases steadily with decreasing time to expiry. This decrease of the dispersion of the estimated PDFs accelerates as it becomes closer to the expiration date. The 50th percentile price is always below the daily NYMEX closing price. Eventually, the futures price converges to the 50th percentile price as the futures contract expires. The figure also conveys information about changes in the skewness and kurtosis of the implied distributions. The location of the mean relative to the lower and upper percentiles is informative of the direction and extent of the skewness and kurtosis. These results confirm similar findings for other futures markets such as currency futures (Bahra 1997). Figure 8.11 shows the sample implied PDFs at four selected dates. It clearly shows that the dispersion of the estimated PDFs around the mean decreases significantly from 27 June 2007 to 24 August 2007, indicating the resolution of uncertainty (about the settlement price for the September 2007 futures contract) with time. The shapes of the estimated PDFs become more symmetric closer to expiration.

TABLE 8.5 Comparison of Root-Mean-Square Errors (\$/MMBTu) for Four Different Models. The PDFs are Sampled at 104 Discrete Asset Values. The Daily Option Prices Used are at Closing from 25 May 2007 through 27 August 2007 for the September 2007 NYMEX Natural Gas Futures

INVSP	BS	JD	MLN
0.0027	0.0070	0.0060	0.0059




FIGURE 8.10 Different percentiles of the estimated PDFs as a function of time. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from five strikes in the money to five strikes out of the money. The option expires on 28 August 2007.




FIGURE 8.11 Sample implied PDFs at selected dates. The PDFs are sampled at 104 discrete asset values. The option prices are observed for the strikes from five strikes in the money to five strikes out of the money. The options expire on 28 August 2007.

8.5 Conclusions

In this paper, we have formulated a general framework for applying the classic inverse theory to the estimation of the implied risk-neutral PDF from option prices. The proposed inversion technique is applied to the inversion of option prices associated with the September 2007 NYMEX natural gas futures and is compared with the popular estimation models of Black–Scholes, a mixture of two lognormals, jump diffusion and the Edgeworth expansion. It is shown that although the fit for the inversion model is significantly better, the fit for the other models is also reasonably acceptable. However, a reasonable fit is not sufficient for a meaningful solution. It is common that the model space is defined by a few parameters, which are then obtained from data by standard least squares. A solution is found by having a reasonable fit to the data. In this kind of fitting procedure, the choice of model space and its scarce parameterization may not be appropriate in terms of the physics of the particular problem and/or recovering the true risk-neutral PDF. Constraining information, which is implied by this choice, is either hidden behind the method or is only justified as convenient, and gives a very narrow class of admissible solutions.

The inverse technique is an integrated approach that is based on both mathematical and physical principles appropriate for the particular problem. Since estimating the implied risk-neutral PDF from finite observed option prices is an ill-posed inverse problem, the basic mathematical questions of existence, uniqueness and stability of the solution have to be considered first. In these considerations, it is necessary to utilize constraining information that comes from the physics of the problem. In this paper, *a priori* information such as the smoothness requirement concerning the shape of the estimated PDF has been introduced through a penalty function. The positivity constraints are included as a hard bond on the PDF values. The best solution is not the solution that gives the best fit, but the solution that gives the optimal trade-off between the goodness-of-fit and the smoothness of the estimated PDF. It is shown that the inversion model not only gives the best goodness-of-fit, but also a significantly better model resolution than the other models. The BS, JD and MLN models basically cannot resolve the densities far away from the strikes where option prices are observed but can resolve long-wavelength features of the densities inside the strikes where option prices are observed. On the other hand, the inversion model can resolve not only the significant details of the densities inside the strikes where option prices are observed, but also the long-wavelength features of the densities away from the strikes where option prices are observed.

The empirical study for the last 3 months of the September 2007 futures contract shows that the shapes of the estimated PDFs become more symmetric as the futures contract becomes closer to the expiration date. The dispersion of the estimated PDFs decreases with decreasing time to expiry, indicating the resolution of uncertainty with passing time.

Acknowledgements

We would like to thank three anonymous referees for their critical and constructive reviews of the manuscript. Yijun Du and Christine Wang thank Dan Kuehn and Jim Gleitman for their support throughout the course of the project.

References

- Abken, P.A., Madan, D. and Ramamurtie, S., Estimation of risk neutral densities by Hermite polynomial approximation: with an application to Eurodollar future options. Federal Reserve Bank of Atlanta, Working Paper 96-5, 1996.
- Ait-Sahalia, Y. and Lo, A.W., Nonparametric estimation of state-price densities implicit in financial asset prices. *J. Financ.*, 1998, **53**, 499–547.
- Avellaneda, M., Friedman, C., Holmes, R. and Samperi, D., Calibrating volatility surfaces via relative-entropy minimization. *Appl. Math. Financ.*, 1997, **4**, 37–64.
- Backus, G.E., Inference from inadequate and inaccurate data, I. *Proc. Natl Acad. Sci. USA.*, 1970, **65**, 1–105.
- Backus, G.E., Comparing hard and soft prior bounds in geophysical inverse problems. *Geophys. J.*, 1988, **94**, 249–261.
- Backus, G.E. and Gilbert, J.F., Numerical applications of a formalism for geophysical inverse problems. *Geophys. J. R. Astr. Soc.*, 1967, **13**, 247–276.
- Backus, G.E. and Gilbert, J.F., Uniqueness in the inversion of inaccurate gross earth data. *Philos. Trans. R. Soc. Lond.*, 1970, **266**, 123–192.
- Bahra, B., Implied risk-neutral probability density functions from options prices: theory and application. Bank of England Working Paper Series, 66, 1997.
- Bates, D.S., Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. *Rev. Financ. Stud.*, 1996, **9**, 69–107.
- Baumeister, J., *Stable Solution of Inverse Problems*, 1987 (Vieweg & Sohn: Friedr).
- Bennett, A.F., *Inverse Modeling of the Ocean and Atmosphere*, 2002 (Cambridge University Press: Cambridge).

- Berestycki, H., Busca, J. and Florent, I., An inverse parabolic problem arising in finance. *C. R. Acad. Sci. Paris Ser. I Math.*, 2000, **331**, 965–969.
- Berestycki, H., Busca, J. and Florent, I., Asymptotics and calibration of local volatility models. *Quant. Finance*, 2002, **2**, 61–69.
- Berestycki, H., Busca, J. and Florent, I., Computing the implied volatility in stochastic volatility models. *Commun. Pure Appl. Math.*, 2004, **57**, 1352–1373.
- Black, F. and Scholes, M., The pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**, 637–654.
- Bouchouev, I. and Isakov, V., The inverse problem of option pricing. *Inverse Probl.*, 1997, **13**, L11–L17.
- Bouchouev, I. and Isakov, V., Uniqueness, stability and numerical methods for inverse problems that arises in financial markets. *Inverse Probl.*, 1999, **15**, R95–R116.
- Breedon, D. and Litzenberger, R.H., Prices of state-contingent claims implicit in option prices. *J. Bus.*, 1978, **51**, 621–651.
- Buchen, P. and Kelly, M., The maximum entropy distribution of an asset inferred from option prices. *J. Financ. Quant. Anal.*, 1996, **31**, 143–159.
- Campa, J.M., Chang, P.H. and Reider, R.L., Implied exchange rate distributions: evidence from OTC option markets. *J. Int. Money Financ.*, 1998, **17**, 117–160.
- Coleman, T., Li, Y. and Verma, A., Reconstructing the unknown local volatility function. *J. Comput. Financ.*, 1999, **2**, 77–102.
- Cont, R., Beyond implied volatility: extracting information from options prices. In *Econophysics*, edited by J. Kertesz and I. Konkor, 1997 (Dordrecht: Kluwer).
- Cont, R. and Tankov, P., *Financial Modeling with Jump Processes*, 2004a (Chapman & Hall/CRC Press: London).
- Cont, R. and Tankov, P., Non-parametric calibration of jump-diffusion option pricing models. *J. Comput. Financ.*, 2004b, **7**, 1–49.
- Cont, R. and Tankov, P., Retrieving Lévy processes from option prices: regularization of an ill-posed inverse problem. *SIAM J. Control Optim.*, 2006, **45**, 1–25.
- Corrado, C.J. and Su, T., Skewness and kurtosis in S&P 500 index returns implied by option prices. *J. Financ. Res.*, 1996, **19**, 175–192.
- Cox, J.C. and Ross, S.A., The valuation of options for alternative stochastic processes. *J. Financ. Econ.*, 1976, **3**, 145–166.
- Crépey, S., Calibration of the local volatility in a generalized Black-Scholes model using Tikhonov regularization. *SIAM J. Math. Anal.*, 2003a, **34**, 1183–1206.
- Crépey, S., Calibration of the local volatility in a trinomial tree using Tikhonov regularization. *Inverse Probl.*, 2003b, **19**, 91–127.
- Du, Y., Aydin, A. and Segall, P., Comparison of various inversion techniques as applied to the determination of a geophysical deformation model for the 1983 Borah Peak earthquake. *Bull. Seismol. Soc. Am.*, 1992, **82**, 1840–1866.
- Engl, H.W., Calibration problems – an inverse problems view. *WILMOTT Mag.*, 2007, July, 16–20.
- Engl, H.W., Hanke, M. and Neubauer, A., *Regularization of Inverse Problems*, 1996 (Kluwer Academic: Dordrecht).
- Franklin, J.N., Well-posed stochastic extension of ill-posed problems. *J. Math. Anal. Appl.*, 1970, **31**, 682–716.
- Hadamard, J., *Lectures on Cauchy's Problem*, 1923 (Yale University Press: New Haven).
- Hansen, P.C., Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Rev.*, 1992, **34**, 561–580.
- Hansen, P.C., Rank-Deficient and Discrete Ill-Posed Problems, 1998 (SIAM: Philadelphia).
- Heston, S.L., A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.*, 1993, **6**, 327–345.
- Hofmann, B., *Regularization for Applied Inverse and Ill-Posed Problems*, 1986 (Leipzig: Teubner).
- Jackson, D.D., Interpretation of inaccurate, insufficient and inconsistent data. *Geophys. J. R. Astr. Soc.*, 1972, **28**, 97–109.

- Jackson, D.D., The use of *a priori* data to resolve non-uniqueness in linear inversion. *Geophys. J. R. Astr. Soc.*, 1979, **57**, 137–157.
- Jackwerth, J.C., Option-implied risk-neutral distributions and implied binomial trees: a literature review. *J. Deriv.*, 1999, **7**, 66–82.
- Jackwerth, J.C. and Rubinstein, M., Recovering probability distributions from option prices. *J. Financ.*, 1996, **51**, 1611–1631.
- Jarrow, R.A. and Rudd, A., Approximate option valuation for arbitrary stochastic processes. *J. Financ. Econ.*, 1982, **10**, 347–369.
- Jondeau, E. and Rockinger, M., Reading the smile: the message conveyed by methods which infer risk neutral densities. *J. Int. Money Financ.*, 2000, **19**, 885–915.
- Jondeau, E. and Rockinger, M., Gram–Charlier densities. *J. Econ. Dynam. Control*, 2001, **25**, 1457–1483.
- Kindermann, S. and Pikkarainen, H.K., Regularization of inverse problems and its application to the calibration of option price models. In *Advanced Financial Modelling*, edited by H. Albrecher, W.J. Runggaldier and W. Schachermayer, pp. 223–244, 2009 (Walter de Gruyter: Berlin).
- Lagnado, R. and Osher, S., A technique for calibrating derivative security pricing models: numerical solution of the inverse problem. *J. Comput. Financ.*, 1997, **1**, 13–25.
- Lanczos, C., *Linear Differential Operators*, 1961 (D. Van Nostrand: New York).
- Lawson, C.L. and Hanson, R.J., *Solving Least Squares Problems*, 1974 (Prentice-Hall: Englewood Cliffs, NJ).
- Levenberg, K., A method for the solution of certain non-linear problems in least-squares. *Q. Appl. Math.*, 1944, **2**, 162–168.
- Longstaff, F., Option pricing and the martingale restriction. *Rev. Financ. Stud.*, 1995, **8**, 1091–1124.
- Madan, D.B. and Milne, F., Contingent claims valued and hedged by pricing and investing in a basis. *Math. Finance*, 1994, **4**, 223–245.
- Malz, A.M., Using options prices to estimate realignment probabilities in European monetary system. Federal Reserve Bank of New York Staff Reports, No. 5, 1995.
- Malz, A.M., Estimating the probability distribution of future exchange rates from option prices. *J. Deriv.*, 1997, **5**, 20–36.
- Marquardt, D.W., An algorithm for least squares estimation of nonlinear parameters. *SIAM J.*, 1963, **11**, 431–441.
- Melick, W.R. and Thomas, C.P., Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf crisis. *J. Financ. Quant. Anal.*, 1997, **32**, 91–115.
- Menke, W., *Geophysical Data Analysis: Discrete Inverse Theory*, 1984 (Academic Press: Orlando, FL).
- Morozov, V.A., *Methods for Solving Incorrectly Posed Problems*, 1984 (Springer: New York).
- Nashed, M.Z., Inner, outer, and generalized inverses in Banach and Hilbert spaces. *Numer. Funct. Anal. Optimiz.*, 1987, **9**, 261–325.
- Natenberg, S., *Option Volatility & Pricing: Advanced Trading Strategies and Techniques*, 1994 (McGraw-Hill: New York).
- Natterer, F., *The Mathematics of Computerized Tomography*, 2001 (SIAM: Philadelphia).
- Oviedo, R., The suboptimality of early exercise of futures-style options: a model-free result, robust to market imperfections and performance bond requirements. Working paper, Social Science Research Network (SSRN), 2006. Available online at: <http://ssrn.com/>
- Parker, R.L., Understanding inverse theory. *Annu. Rev. Earth Planet. Sci.*, 1977, **5**, 35–64.
- Ritchey, R.J., Call option valuation for discrete normal mixtures. *J. Financ. Res.*, 1990, **13**, 285–296.
- Rubinstein, M., Implied binomial trees. *J. Financ.*, 1994, **49**, 771–818.
- Rubinstein, M., Edgeworth binomial trees. *J. Deriv.*, 1998, **5**, 20–27.
- Sherrick, B.J., Garcia, P. and Tirupattur, V., Recovering probabilistic information from option markets: tests of distributional assumptions. *J. Fut. Mkts*, 1996, **16**, 545–560.
- Shimko, D., Bounds of probability. *Risk*, 1993, **6**(4), 34–37.
- Stutzer, M., A simple nonparametric approach to derivative security valuation. *J. Financ.*, 1996, **51**, 1633–1652.

- Tarantola, A., *Inverse Problem Theory, Methods for Data Fitting and Model Parameter Estimation*, 1987 (Elsevier: Amsterdam).
- Tarantola, A. and Valette, B., Generalized nonlinear inverse problems solved using the least squares criterion. *Rev. Geophys. Space Phys.*, 1982, **20**, 219–232.
- Tikhonov, A.N. and Arsenin, V.Y., *Solutions of Ill-Posed Problems*, 1977 (Wiley: New York).
- Wiggins, R.A., The general linear inverse problems: implication of surface waves and free oscillations for earth structure. *Rev. Geophys. Space Phys.*, 1972, **10**, 251–285.

Investing in the Wine Market: A Country-Level Threshold Cointegration Approach

Lucia Baldi

Massimo Peri

Daniela Vandone

9.1	Introduction	173
9.2	Wine: The Theoretical Framework.....	175
9.3	The Mediobanca Index: Evidence on Stock Performance.....	176
	The Data • Performance	
9.4	Econometric Methodology.....	180
9.5	Empirical Results.....	181
9.6	Conclusion	186
	Reference	186

9.1 Introduction

In recent years, investments in commodities, via commodity futures and commodity index funds, have grown rapidly. The appeal of investing in commodities is generally attributed to the low correlation with traditional stocks, which allows for portfolio diversification benefits, that is reduction of risk for any given level of expected return (Erb and Harvey 2006, Gorton and Rouwenhorst 2006, Sanning et al. 2007, Geman and Kharoubi 2008, Masset and Henderson 2008, Buyuksahin et al. 2010, Chong and Miffre 2010, Masset and Weisskopf 2010), and the evidence that profitable trading strategies can be constructed. Indeed, a number of studies find that commodity returns can be predicted by a number of variables (Hong and Yogo 2011), and profitable momentum and term structure strategies can be implemented (Schwarz and Szakmary 1994, Crain and Lee 1996, Brooks et al. 2001, Yang et al. 2001, Erb and Harvey 2006, Gordon et al. 2007, Miffre and Rallis 2007, Fuertes et al. 2010, Stoll and Whaley 2010, Szakmary et al. 2010). The majority of these studies were conducted in the framework of cointegration analysis and threshold cointegration analysis; from a methodological standpoint, this allows the analysis of the interdependence between asset classes, their relationships, and the speed of adjustment to the long-run equilibrium.

In our paper we focus on wine, a commodity of growing importance not only in European countries (often referred to as ‘Old World Countries’), that have historically dominated this market, but also in the so-called ‘New World Countries’, e.g. China, South Africa, Chile, that have registered astonishing rates of growth, at times with extraordinary speed, entering not only the lower-quality segment, but also the medium–high segment, once the exclusive domain of traditional long-established producers (Aylward 2003, Aylward and Turpin 2003).

The aim of this study is to analyse the long-run relationship between wine share price indexes and general stock market indexes, examining their different speeds of adjustment to the long-run equilibrium. The goal is to provide traders with signals of information inefficiency that could be exploited to make profitable investment strategies.

Indeed, according to Fama (1965), in an efficient market all available information is instantaneously and completely reflected in stock prices. Thus, it is not possible to forecast future price evolution on the basis of previous stock price variations because this information is already integrated in the present price. Conversely, if there is no efficiency in the market, then active traders can anticipate price movements over the short run and make profitable investments from buying and selling stocks (Siourounis 2002). Specifically, if two economic series, such as financial indexes, are cointegrated, this implies that movements of one series can be used to predict fluctuations of the other: in other words, traders can anticipate the evolution of the dynamics of a stock market index by knowing that of another stock market index: the key question we pose is whether investors can exploit the dynamics of stock markets to predict wine indexes returns and thus make profitable investments by buying and selling stocks.

Following the seminal paper of Balke and Fomby (1997), we use threshold cointegration rather than cointegration since this econometric technique allows for non-linear adjustments in the long-run equilibrium (Hokkio and Rush 1989, Anderson 1997, Perez-Quiros and Timmermann 2000, Maasoumi and Racine 2002). Indeed, since the efficiency hypothesis assumes no transaction costs, free and symmetric information, as well as rational investors, studying stock prices and efficiency using linear cointegration techniques corresponds to the assumption of symmetric, linear and continuous stock price adjustment dynamics. These set of assumptions seem to be very constraining, since markets present friction, such as transaction costs, noise traders and imitative behaviour, and this can imply price adjustment dynamics towards fundamental values that are discontinuous and non-linear (Enders and Siklos 2001, Shively 2003, Prat and Jawadi 2009).

The economic rationale for considering the possibility of a non-linear rather than a linear type of adjustment to the long-run equilibrium is that the first allows prices to adjust in a different way to large or small deviations from the long-run equilibrium level. This implies that the dynamic behaviour of the rate of return differs according to the size of the deviation. In fact, this methodology captures those adjustments that are active only when deviations from the equilibrium exceed a threshold, which is often represented by transaction costs (Sercu et al. 1995, Jawadi and Koubaa 2004, Aslandis and Kouretas 2005). In fact, traders may not act immediately as prices move, due to the possibility of 'mis-price deepening' (Shleifer 2000), they only act when the expected profits exceed the costs. In this sense, the threshold cointegration constitutes an adequate specification, since the error correction mechanism is active only above a certain size of the variation compared with the equilibrium.

In other words, in the linear cointegration approach the adjustment parameters are assumed to be constant within the period analysed, while in the threshold cointegration approach the error correction terms (ECTs) are inactive when the value is inside a given range, but are active above a certain threshold. When the deviation from equilibrium is above the critical threshold, agents enter the market to implement profitable arbitrage activities, moving the system back to the equilibrium (McMillan 2003, 2005).

This study makes use of the Mediobanca Global Wine Industry Share Price Index and the composite stock market indexes for a selection of countries belonging to the Old World Countries and the New World Countries—Australia, Chile, China, France and the United States—where wine shares are listed. The Mediobanca Wine Index encompass companies operating in the wine industry listed on stock markets worldwide.

The originality of this work is twofold. Firstly, the threshold cointegration methodology is applied to the wine sector; to the authors' knowledge, no previous study has investigated the long-term dynamics between the share price of wine companies and whole stock market indexes. However, as underlined above, such an analysis is relevant since it enables a better understanding of financial investment opportunities due to wine share price adjustment dynamics and stock market inefficiencies. Secondly,

the focus of the analysis is not on fine wine, as in most of the current literature, but on non-fine wine (which will be referred to as normal wine in the remaining part of the article). The interest in normal wine stocks as an asset class that traders can use for investment purposes is due to the growing size of the wine market in countries not historically suited to the production of wine (such as China, Australia and Chile, where large companies have recently become established in the market), their increasing stock price volatility and better performance. This paper presents the first use of the Mediobanca Global Wine Industry Share Price Index in an academic context, presenting this databank to the wider research community through one of the uses that can be made of this resource.

The paper is organized as follows. Section 9.2 describes the theoretical framework. Section 9.3 presents the dataset used for the purpose of the study and a brief analysis of index performance. Section 9.4 proposes the econometric methodology and Section 9.5 develops the empirical results. Section 9.6 includes the discussion and final remarks.

9.2 Wine: The Theoretical Framework

The economic literature on commodities presents a wide range of interesting studies that have analysed the potential of wine as a financial investment (Fogarty (2006) presents a detailed overview of the subject). Most investigated has been the estimation of the rate of return of wines, comparing this value with the return from other common assets, both in the mean value and in the conditional volatility or covariance (see, among others, Krasker (1979), Weil (1993), Burton and Jacobsen (2001) and Sanning *et al.* (2007)).

This literature has been interested only in fine and rare wines, and, using auction data of specific wines or composite indexes (e.g. *LiveEx*), its results have encompassed a broad mix of empirical findings that has provided financial guidance in areas such as prices (Jaeger 1981, Weil 1993, Fogarty 2006), buyer's premium (De Vittorio and Ginsburgh 1996), speculative bubbles (Jovanovic 2007), excellence of vintage and respective ranking (Jones and Storchmann 2001, Masset and Henderson 2008), fluctuations in inventories (Bukanya and Labys 2007) and portfolio diversification strategies. Specifically, as far as diversification is concerned, Sanning *et al.* (2007) use the Fama–French three-factor model and the Capital Asset Pricing Model to directly assess the risk–return profile of wines as compared with equities and find that investment-grade wines benefit from low exposure to market risk factors, thus offering a valuable dimension for portfolio diversification. Masset and Henderson (2008) find that investing in the wine market might achieve an attractive performance in terms of both average returns and volatility since wine returns are only slightly correlated with other assets and, as such, they can be used to reduce the risk of an equity portfolio. Moreover, Fogarty (2006) finds that the performance of Australian wines is comparable to that of Australian equities, and Burton and Jacobsen (2001) demonstrate that the returns of a wine portfolio consisting only of wines from the 1982 vintage compare favourably with that of the Dow Jones.

Among these studies, French fine wines are the most commonly analysed, together with Australian high-quality production.

Despite the extensive empirical literature on the wine market, to the best of our knowledge there are no studies that have analysed, from a financial viewpoint, non-fine wine. However, in recent years the importance of normal wine has grown rapidly, specifically in New World Countries. Indeed, while historically the wine market has been dominated by European countries (often referred to as Old World countries), since the beginning of the 1990s new producing countries have found their way into the market, showing strong competitive potential due to their innovative strategies in production and trade (Campbell and Guibert 2006). Europe (in particular, France, Italy, Germany and Spain) still occupies a leading position on the world wine market, accounting globally for 49% of growing areas and 63% of wine production (data from the FAO databank for the year 2007). However, wine is also currently produced in Argentina (accounting for 9% of world production), the USA (8%), China (5%), Australia (4%), South Africa (4%) and Chile (3%). Contrary to many traditional winegrowing countries, where

production has dropped by 25% compared with volumes in the 1990s, 'New World' countries have registered astonishing rates of growth, at times with extraordinary speed (as in the case of China), entering not only the lower quality segment, but also the medium-high segment, once the exclusive domain of traditional long-established producers (Aylward 2003, Aylward and Turpin 2003).

9.3 The Mediobanca Index: Evidence on Stock Performance

9.3.1 The Data

The analysis presented in this paper makes use of data on the wine share price index and the composite stock market index of the stock exchange for five countries belonging both to Old and New World Countries: Australia, Chile, China, France and the US. The dataset covers the period starting January 1, 2001 up to the end of February 2009. All series are expressed in euro and appear in the econometric model in logarithmic form. The wine series is the Mediobanca Global Wine Industry Share Price Index from Mediobanca,* which covers companies operating in the wine industry, listed on regulated stock markets and quoted for at least six months. Prices are computed daily and represent a financial benchmark of wine, measuring and monitoring the dynamic of risk and return of wine stocks. The index is calculated for each of those countries whose stock had traded at least three titles that meet some specific selection criteria.†

Data on the composite stock market series are daily prices supplied by Datastream, and represent the performance of the whole stock market for a given country. Specifically, the data used is for the Australian S&P/ASX300 index, the Chilean IPSA index, the Chinese SSE index, the French CAC40 index and the US S&P500 index.‡

All indexes are 'capitalization-weighted', that is the components are weighted according to the total market value of their outstanding shares.

Both series are 'price' indexes, expressing the dynamics of stock prices alone, without the component of income represented by the distribution of dividends. 'Total return' indexes, which also include dividends, are available for all series, but the pure price index is preferred. The rationale behind this choice is that the dividend policy adopted by each company is not relevant in the analysis presented here, as, in general, dividends do not reveal the level of volatility that would be necessary to influence the null hypothesis of 'no cointegration' among a set of share price indices (Dwyer and Wallace 1992, Subramanian 2009).

9.3.2 Performance

Figure 9.1 shows the cumulative stock return for wine indexes over the period from 2001 to 2009. At the end of this eight-year period France, the US and Australia had achieved almost similar total cumulative returns (respectively 37%, 37% and 40%), although with very different dynamics. In particular, the French wine index suffered a decline between the beginning of 2001 and mid-2002, followed by a stable rise and then steeper growth from the beginning of 2005 up to the beginning of 2008. During this last period the French wine index almost tripled. Since mid-2008, however, the index has decreased in line with other financial assets and stock markets as a result of the global economic crisis.

* <http://www.mbres.it/>

† The Mediobanca index includes wine companies selected according to the following characteristics: companies listed on regulated markets; series of quotes of at least six months; at least 50% of revenues must come from initiating wine; commitment as direct management in the production cycle. The panel index is comprised of 42 stocks and has an aggregate market capitalization of €14.3 bn.

‡ The S&P/ASX300 index is the stock market index of Australian stocks listed in the Australian Securities Exchange; the IPSA is a stock market index composed of the 40 stocks with the highest average annual trading volume in the Santiago Stock Exchange; the SSE is the composite index from the Shanghai Stock Exchange; CAC40 is the benchmark French stock market index and includes the 40 most significant stocks in terms of liquidity; and, finally, the S&P500 includes the prices of 500 large-cap common stocks actively traded on the two largest American stock markets, NYSE and NASDAQ.




FIGURE 9.1 Cumulative stock return for wine indexes for the five countries considered.

Source: Our calculations on Mediobanca data.

In Australia, after an initial bearish trend, the wine index showed a stable pattern over the period, while in the US the wine index rose between 2003 and 2006 and later declined, although to a lesser extent than its French counterpart.

Chile and China reached a higher level of cumulative returns over the 2001–2009 period (respectively 101% and 94%), but again with different patterns. In particular, the value range of the Chilean index went from a minimum of 99.50 points to a maximum of 290 points, while in China the index declined significantly until the end of 2005 and then climbed sharply at the beginning of 2008. The strong performance of the index for China continued into 2009.

Estimates of the average daily return and volatility for the wine and composite index were also calculated and are summarized in Table 9.1. For France, the US, Chile and China the average daily return of wine indexes is higher than that of composite indexes (respectively CAC40, S&P500, IPSA, and SSE) over the 2001–2009 period. More specifically, while the average return of wine indexes is always positive, the average return of composite indexes is negative for all countries except Chile. The most significant difference is that from the US, where the average daily return of the wine index is 95% while the S&P500 average daily return is –28%. Looking at yearly intervals, it becomes apparent that only rarely do wine indexes yield negative returns, almost always outperforming the composite indexes.

Australia is the only exception. There, apart from the 2001–2002 interval, the wine index average return is lower than that of the S&P300. The Australian wine sector has achieved many successes in recent decades through government measures promoting exports and low taxes. Moreover, the Australian wine industry is characterized by high levels of concentration (four companies accounting for over 75% of production), providing economies of scale in producing value-for-money wines. However, after a planting boom during the mid to late nineties, from 2001 onwards this country has gone through a very difficult period mainly due to strong supply pressure. As a result, a number of listed

TABLE 9.1 Daily Wine and Composite Index Returns and Volatility

	Average daily return (%)		Cond. volatility ^a (%)	
	Wine	Comp.	Wine	Comp.
Australia				
2001–2002	13.81	−1.24	1.02	1.08
2003–2004	−2.49	19.42	0.90	0.78
2005–2006	0.49	23.18	1.21	0.90
2007–2008 ^b	−12.40	−5.05	1.81	2.69
All period	22.40	30.96	1.37	1.27
Chile				
2001–2002	29.02	−10.40	0.90	1.28
2003–2004	11.33	36.38	0.90	1.13
2005–2006	0.90	32.09	0.80	0.95
2007–2008 ^b	17.19	3.71	1.40	1.73
All period	75.54	32.84	1.05	1.28
China				
2001–2002	−4.19	−14.98	4.54	1.70
2003–2004	2.50	−9.74	1.32	1.38
2005–2006	58.92	7.79	1.67	1.44
2007–2008 ^b	21.80	77.40	2.46	2.69
All period	41.14	−28.39	2.03	1.93
France				
2001–2002	−7.09	−25.73	1.28	1.95
2003–2004	5.02	11.06	0.91	1.02
2005–2006	40.07	21.63	0.90	0.80
2007–2008 ^b	22.83	−11.89	1.45	3.07
All period	67.45	−26.53	1.15	1.78
U.S.				
2001–2002	28.15	−14.58	1.68	1.69
2003–2004	31.91	4.95	1.15	1.13
2005–2006	9.74	12.90	1.37	0.81
2007–2008 ^b	−28.03	−13.58	2.51	1.90
All period	95.05	−28.52	1.87	1.36

Source: Own calculations with Mediobanca and Datastream data.

^a All the volatilities were computed via the GARCH(1,1) model except for the USA composite index within the 2003–2004 range where the GARCH(2,1) model was used.

^b The 2007–2008 range also includes 42 observations for 2009.

companies have announced profit downgrades (including Southcorp company, Australia's largest wine producer), perhaps contributing to making wine stocks less of a profitable investment.

Elsewhere, the risk features of both indexes are more heterogeneous. Chile and France show a lower wine index volatility, whereas for Australia, the US and China the volatility of the wine index is slightly higher than that of the composite. It is interesting to note that during the financial crisis in the period 2007–2008, all countries except the US show a conditional volatility of the wine index that is lower than that of the composite index.

Overall, what emerges from the analysis of the returns is that, in general, wine indexes performed well during the last decade and investors could have earned greater returns by investing in these indexes rather than solely investing in the domestic composite indexes, although in some cases they would have been exposed to greater volatility.

The superior performance of the wine index is clearly visible from the graphs shown in Figure 9.2, which represent the cumulative abnormal return of the wine indexes, that is the market-adjusted abnormal return.

The abnormal return is estimated by subtracting the composite stock market index from the return of the wine index. The resulting evidence is quite interesting. Apart from Australia, whose wine industry suffered a crippling financial crisis during the mid-2000s (see above), the cumulative abnormal returns remained significantly positive for almost all years and in all countries.

From Figure 9.2 it can be observed that the US registered the best performance of the countries under analysis. This is not surprising given that the American wine industry leads the group of New World wine producers and is the world's fourth largest wine producer, however remaining a net importer of wine (Canning and Perez 2008, Goodhue *et al.* 2008). The growth of the sector started several decades




FIGURE 9.2 Abnormal return in each of the five countries considered.

Source: Our calculations on Mediobanca data.

ago and continues to this day, not having had any apparent slowdown recorded since its early days. As in other New World wine industries, firm concentration is particularly high, even if the industry has been evolving recently with the proliferation of new wineries (Insel 2008). In addition, the sector benefits from a concentrated and efficient distribution system for products. These factors together help explain the good performance of US wine stocks in the financial markets.

Elsewhere, the abnormal return for France is high, but lower than that of the US. The country, traditionally one of the largest world producers and consumers of wine, has a fragmented industry subject to too many controls (Terblanche *et al.* 2008). Since 2001, French wines have been going through a slowdown period, losing market share both in domestic and in export markets, a situation exacerbated by negative currency effects arising from a strong euro (Hussain *et al.* 2008).

On the other hand, the Chilean wine sector has recorded a notable production and export record during the last decade and many wine firms have developed the competence to be present in an increasingly large number of countries (Giuliani and Bell 2005, Gwynne 2008). Concha Y Toro, the market leader, has become the most successful Chilean exporting company, transforming its business radically and becoming the world's first winery to list on the New York Stock Exchange. It should not be surprising, therefore, that the Chilean wine sector's performance in terms of return should have been positive.

Finally, the Chinese wine industry, although still oriented towards the domestic market (which is still relatively small and is dominated by the medium- and low-quality segment), has made very rapid progress in recent years (Jenster and Cheng 2008, Mitry *et al.* 2009), with the performance of wine in the stock market having always been satisfactory, surpassing that of the Shanghai Stock Exchange.

As a concluding remark, it can therefore be argued that, in general, wine indexes have performed well compared with composite indexes during the last decade, and that investors could have earned greater returns by investing in this market.

9.4 Econometric Methodology

Evidence of cointegration (Granger 1981) among several indexes of stock prices suggests that series have the tendency to move together in the long run even if they experience short-term deviations from their common equilibrium path (Masih and Masih 1997, Patra and Poshakwale 2008). These traditional models assume that the adjustment process to maintain *equilibrium* occurs at every time period. However, many situations, and in particular stabilizations of commodity prices, are often characterized by discrete interventions. In recent years, two main classes of models have been proposed in the literature to characterize this kind of non-linear adjustment process. One class considers Markov-switching Vector Error Correction models, assigning probabilities to the occurrence of different regimes (Hamilton 1989, Krolzig 1997). The second class is based on Tong and Lim's (1980) approach using a Self-exciting Autoregressive Model where the regimes that have occurred in the past and the present are known with certainty, 'certainty' being established using statistical techniques. In this framework, Balke and Fomby (1997) introduced the concept of 'threshold cointegration', a feasible estimation methodology that allows the adjustment process to move differently in separate regimes. They hypothesized that this movement towards a long-run equilibrium may not occur in every time period, but only when the deviation from equilibrium exceeded a critical threshold. Following Balke and Fomby (1997), in this paper we apply a threshold vector error correction model (TVECM), with a threshold effect based on an error correction term. In the case of two regimes, Balke and Fomby (1997) present a TVECM of order $L + 1$ that takes the form

$$\Delta x_t = \begin{cases} A'_1 X_{t-1}(\beta) + u_t, & \text{if } w_{t-1}(\beta) \leq \gamma \text{ regime1} \\ A'_2 X_{t-1}(\beta) + u_t, & \text{if } w_{t-1}(\beta) > \gamma \text{ regime2} \end{cases} \quad (9.1)$$

where

$$X_{t-1}(\beta) = \begin{bmatrix} 1 \\ w_{t-1}(\beta) \\ \Delta x_{t-1} \\ \Delta x_{t-2} \\ \vdots \\ \Delta x_{t-l} \end{bmatrix} \quad (9.2)$$

and x_t is a p -dimensional time series $I(1)$ cointegrated with one $(p \times 1)$ cointegrating vector β , $w_t(\beta)$ is the ECT, and u_t is the error term assumed to be an iid Gaussian sequence with a finite covariance matrix. A_1 and A_2 are matrices of coefficients describing the dynamics in each regime, while γ is the threshold parameter. Values of w_{t-1} below or above the threshold γ allow the coefficients to switch between regimes 1 and 2; in particular, the estimated coefficients of w_{t-1} of each regime denote the different adjustment speeds of the series towards equilibrium.

Hansen and Seo (2002) provided an estimation method for TVECM via maximum likelihood, which involves a joint grid search over the threshold parameter and cointegrating vector. In order to test for threshold cointegration, Tsay (1989, 1998) proposed non-parametric non-linearity tests, while Andrews (1993), Hansen (1996), Balke and Fomby (1997) and Lo and Zivot (2001) presented different methods of estimation based on the Lagrange Multiplier (LM) statistics. More recently, Hansen and Seo (2002) developed two SupLM (Supremum Lagrange Multiplier) tests for a given or estimated β using a parametric bootstrap method to calculate asymptotic critical values with the respective p -values. The first test is denoted as

$$\sup LM^0 = \sup_{\gamma_L \leq \gamma \leq \gamma_U} LM(\beta_0, \gamma)$$

and would be used when the true cointegrating vector β is known *a priori*. The second test is used when the true cointegrating vector $\tilde{\beta}$ is unknown and the test statistic in this more general case corresponds to

$$\sup LM = \sup_{\gamma_L \leq \gamma \leq \gamma_U} LM(\tilde{\beta}, \gamma)$$

where $\tilde{\beta}$ is the null estimate of the cointegrating vector. In these tests, the search region $[\gamma_L, \gamma_U]$ is set so that γ_L is the π_0 percentile of $\tilde{w}_{t-1}$, where $\tilde{w}_{t-1} = w_{t-1}(\tilde{\beta})$, and γ_U is the $(1 - \pi_0)$ percentile.*

9.5 Empirical Results

To implement the asymmetric cointegration approach, we carried out several steps. First, we tested the degree of integration of the variables via the Augmented-Dickey Fuller test (ADF) and the Philips–Perron test (PP). Subsequently, cointegration (Johansen 1988, Johansen and Juselius 1990) and Granger causality (Granger 1969) between the price pairs (wine and composite index) were tested for each of the countries analysed. The following step entailed a test for the presence of threshold cointegration. Finally, TVECM was run using the Hansen and Seo (2002) procedure.

Table 9.2 shows the results of the ADF and PP tests, where Δ in front of variable names indicates the differentiated series. It emerges that all the series are $I(1)$ with and without trend.

Since the price series have a unit root, the presence of cointegration between the series can be tested following the Johansen approach using the Trace and Maximum-Eigenvalue tests. Both tests were

* Andrews (1993) argued that setting π_0 between 0.05 and 0.15 is generally a good choice.

TABLE 9.2 Test for Unit Root and Stationarity^a

		No trend		With trend	
		ADF	PP	ADF	PP
Australia	Wine	-1.904	-1.817	-2.442	-2.331
	Δ Wine	-45.599	-45.665	-45.588	-45.653
	Composite	-0.980	-0.952	0.191	0.294
	Δ Composite	-45.995	-46.007	-46.032	-46.048
Chile	Wine	-2.312	-2.304	-2.160	-2.432
	Δ Wine	-42.741	-43.048	-42.767	-43.063
	Composite	-0.770	-0.895	-1.508	-1.713
	Δ Composite	-38.502	-38.390	-38.493	-38.381
China	Wine	-0.561	-0.561	-1.502	-1.503
	Δ Wine	-44.661	-44.638	-44.659	-44.636
	Composite	-1.098	-1.130	-1.205	-1.230
	Δ Composite	-46.186	-46.211	-46.229	-46.248
France	Wine	-0.803	-0.866	0.887	0.371
	Δ Wine	-43.223	-43.858	-43.248	-43.870
	Composite	-1.360	-1.110	-1.329	-1.060
	Δ Composite	-48.517	-48.975	-48.508	-48.967
U.S.	Wine	-1.869	-1.855	-0.604	-0.522
	Δ Wine	-46.838	-46.846	-46.922	-46.940
	Composite	-1.306	-0.877	-1.734	-1.298
	Δ Composite	-48.517	-48.975	-48.508	-48.967

^a 1% critical value: ADF and PP -3.430; ADF and PP with trend -3.960.

TABLE 9.3 Cointegration Test between the Wine and Composite Indexes

Series	Hypothesized No. of CE(s)	Trace test	0.05 critical value	Max-Eigen test	0.05 critical value
Australia	None	6.882	15.410	6.270	14.070
	At most 1	0.612	3.760	0.612	3.760
Chile	None	12.227	15.410	10.363	14.070
	At most 1	1.864	3.760	1.864	3.760
China	None	8.124	15.410	5.730	14.070
	At most 1	2.394	3.760	2.394	3.760
France	None	23.735	15.410	22.640	14.070
	At most 1	1.097	3.760	1.097	3.760
U.S.	None	6.591	15.410	6.290	14.070
	At most 1	0.301	3.760	0.301	3.760

Lag(s) interval: lag=1 for Australia and China, lag=2 for France, United States and Chile (selected by the Akaike Information Criterion in VAR). Trend assumption: linear deterministic trend.

conducted including an intercept in the cointegrating equations and estimating the model with a linear trend. The results shown in Table 9.3 indicate the presence of a linear cointegration relationship only in France. In the other countries the results indicate the absence of a cointegration vector at the 0.05 critical value, leading to the conclusion that the Australian, Chilean, Chinese and American wine share price indexes and composite stock market indexes have an unlikely long-term linear relationship.

Turning to market information efficiency, cointegration estimates may lead to the conclusion that, generally, traders cannot use movements of one series to predict fluctuations of the other.

However, as already pointed out, traders act on the market only when expected profits exceed costs, such as transaction fees and other market friction, that is only when the deviation from the long-run equilibrium is over a critical threshold. Consequently, we perform a threshold cointegration analysis. The presence of a threshold was estimated via the application of the Hansen and Seo (2002) SupLM test (when β is estimated) using a parametric bootstrap method with 3000 replications. The results of the tests are reported in Table 9.4. The residual bootstrap value of the SupLM test provides evidence of the presence of threshold cointegration for all the cases studied except for Australia.

The robustness of the threshold results is also supported by the rejection of the null of the equality of ECT coefficients between the two regimes detected. Apart from the Australian case of no threshold cointegration, the p -value of the Wald test is significant for all countries.

The threshold value identified in the French series detects the presence of two regimes with different adjustment speeds in the long-run equilibrium. Table 9.5 reports the estimated coefficients for the TVECMs and the related graphs that exhibit the error correction effect, i.e. the estimated regression functions of the wine and composite index as a function of ECT, holding the other variables constant. The first regime, defined as the *usual regime*, includes the majority (94%) of the observations, while the second, defined as *unusual*, contains the remaining 6% of observations. As can be seen from the figure,

TABLE 9.4 Threshold Cointegration Test

	Australia	Chile	China	France	USA
Test statistic value (SupLM)	15.177	25.282	20.650	23.178	24.500
Residual bootstrap value	0.133	0.049	0.030	0.066	0.040
Threshold value		3.480	-1.348	-9.477	0.284
Estimate of the cointegration vector		0.381	1.562	3.624	1.212
Wald test for equality of ECM coefficient		5.416	10.736	8.231	7.242
p -Value		0.067	0.005	0.016	0.027

TABLE 9.5 Threshold VECMs between the French Wine Index and the CAC40 Composite Index

	Usual regime		Unusual regime	
	(94% of obs.)		(6% of obs.)	
	Wine	Comp.	Wine	Comp.
W_{t-1}	-0.001 <i>-1.830</i>	0.000 <i>0.696</i>	-0.008 <i>-0.520</i>	0.057 <i>2.480</i>
Intercept	-0.009 <i>-1.730</i>	0.005 <i>0.639</i>	-0.074 <i>-0.538</i>	0.533 <i>2.466</i>
ΔWine_{t-1}	-0.009 <i>-0.315</i>	-0.013 <i>-0.376</i>	0.037 <i>0.330</i>	-0.143 <i>-0.757</i>
ΔWine_{t-2}	0.000 <i>-0.012</i>	0.014 <i>0.347</i>	0.273 <i>2.070</i>	0.166 <i>0.768</i>
$\Delta \text{Composite}_{t-1}$	0.080 <i>3.541</i>	-0.076 <i>-2.504</i>	0.160 <i>2.665</i>	0.143 <i>1.455</i>
$\Delta \text{Composite}_{t-2}$	0.037 <i>1.795</i>	-0.016 <i>-0.466</i>	0.020 <i>0.263</i>	-0.090 <i>-0.946</i>

Note: t -Student obtained by Eicker-White standard errors and reported in italics.

in the usual regime the ECT coefficients are quite close to zero, indicating that the variables are close to a random walk.

In the unusual regime the speed of the adjustment coefficient of the composite index (CAC40) is significant and higher with respect to the wine index. In particular, when the gap between the two price series exceeds a critical threshold ($\gamma > -9.477$) the speed of the domestic stock market index's response in restoring the long-run equilibrium is seven time faster than the wine share price index. Therefore, assuming that the long-run relationship is governed by the composite index, the different speeds of adjustment could be used by investors to achieve profitable gains. Hence, when the price gap is over a critical threshold, informed investors operating in the wine sector exploiting market inefficiency can make gainful investments just by looking at the price adjustment dynamics of the domestic stock market and forecasting fluctuations in wine stocks.*

Considerations similar to those relating to the French case could be formulated for the United States, whose TVECM results are reports in Table 9.6. Here the usual regime includes 61% of observations and, as for France, it is close to a random walk. Only when the price gap is over a critical threshold does the adjustment coefficient become active in restoring the long-run equilibrium. In contrast to the findings for France, however, in the United States the speed of adjustment of the composite index to the long-run equilibrium is slower, but the adjustment process involves substantially more observations (39%). On the other hand, the composite index is three times faster than the wine share price index's response to the disequilibrium. Hence, also in the United States there is a boundary of profitable arbitrage in managing wine share price indexes. In particular, in the United States the differences in the adjustment speed are smaller, but more frequent than in France. Consequently, agents have more opportunities to make profitable investments, but with a shorter operating time span.

The findings for France and the United States differ from those related to the two developing countries analysed here, whose TVECM results appear in Tables 9.7 and 9.8. In Chile, the unusual regime

TABLE 9.6 Threshold VECMs between the North America Wine Index and the S&P500 Composite Index

	Usual regime		Unusual regime	
	(61% of obs.)		(39% of obs.)	
	Wine	Comp.	Wine	Comp.
W_{t-1}	-0.001 <i>-0.503</i>	0.001 <i>1.039</i>	0.006 <i>0.790</i>	0.018 <i>2.811</i>
Intercept	0.001 <i>1.821</i>	0.000 <i>-0.635</i>	-0.004 <i>-1.176</i>	-0.008 <i>-3.059</i>
ΔWine_{t-1}	-0.005 <i>-0.113</i>	0.000 <i>-0.008</i>	0.015 <i>0.308</i>	0.005 <i>0.111</i>
ΔWine_{t-2}	0.019 <i>0.530</i>	-0.004 <i>-0.108</i>	0.033 <i>0.687</i>	0.006 <i>0.159</i>
$\Delta \text{Composite}_{t-1}$	-0.063 <i>-1.804</i>	-0.103 <i>-2.824</i>	-0.082 <i>-1.235</i>	-0.214 <i>-3.886</i>
$\Delta \text{Composite}_{t-2}$	0.029 <i>0.665</i>	-0.026 <i>-0.656</i>	-0.250 <i>-2.912</i>	-0.115 <i>-1.320</i>

Note: *t*-Student obtained by Eicker-White standard errors and reported in italics.

* For the French case, all these results may appear ambiguous. In fact, the French series results are cointegrated but also threshold cointegrated. Moreover, we obtained a high level of threshold values and a small percentage of observations in the unusual regime. These results may lead to the conclusion that, in this case, the model just identifies outliers rather than genuine non-linearity behaviour. However, we obtained good statistic results for the significance of the threshold, allowing rejection of the null of the equality of the ECT coefficients.

TABLE 9.7 Threshold VECMs between the Chile Wine Index and the IPSA Composite Index

	Usual regime		Unusual regime	
	(90% of obs.)		(10% of obs.)	
	Wine	Comp.	Wine	Comp.
W_{t-1}	-0.002	0.003	-0.056	-0.064
	<i>-1.151</i>	<i>1.511</i>	<i>-1.723</i>	<i>-1.933</i>
Intercept	0.007	-0.010	0.193	0.225
	<i>1.253</i>	<i>-1.501</i>	<i>1.705</i>	<i>1.941</i>
ΔWine_{t-1}	0.028	0.013	0.156	-0.096
	<i>0.858</i>	<i>0.330</i>	<i>1.642</i>	<i>-1.196</i>
ΔWine_{t-2}	0.069	0.100	0.059	0.087
	<i>2.170</i>	<i>2.378</i>	<i>0.561</i>	<i>1.258</i>
$\Delta \text{Composite}_{t-1}$	0.064	0.187	0.193	0.224
	<i>2.942</i>	<i>5.319</i>	<i>1.695</i>	<i>3.390</i>
$\Delta \text{Composite}_{t-2}$	0.017	-0.080	-0.300	-0.144
	<i>0.755</i>	<i>-2.018</i>	<i>-3.003</i>	<i>-1.593</i>

Note: t -Student obtained by Eicker-White standard errors and reported in italics.

TABLE 9.8 Threshold VECMs between the China Wine Index and the SSE Composite Index

	Usual regime		Unusual regime	
	(68% of obs.)		(32% of obs.)	
	Wine	Comp.	Wine	Comp.
W_{t-1}	0.002	0.002	0.017	0.016
	<i>1.305</i>	<i>1.258</i>	<i>3.490</i>	<i>3.270</i>
$\Delta \text{Intercept}$	0.004	0.004	0.018	0.018
	<i>1.319</i>	<i>1.055</i>	<i>3.509</i>	<i>3.539</i>
ΔWine_{t-1}	-0.048	-0.009	0.198	0.084
	<i>-1.189</i>	<i>-0.200</i>	<i>3.226</i>	<i>1.460</i>
$\Delta \text{Composite}_{t-1}$	0.042	0.025	-0.169	-0.100
	<i>0.909</i>	<i>0.562</i>	<i>-2.997</i>	<i>-1.812</i>

Note: t -Student obtained by Eicker-White standard errors and reported in italics.

starts as a critical threshold of 3.48, but this includes only 10% of observations, while in China, where the threshold is -1.35, the usual regime includes 39% of observations. In both countries the ECTs of the usual regime exhibit small levels of significance and minimal dynamics, whilst becoming significant for both indexes in the unusual regime.

In the case of Chile, the coefficient of the speed of adjustment of the composite index is slightly larger than for the wine coefficient, providing evidence of a reduced profitable space for investment.

In China, where most of the shares are held by Chinese retail investors, the results need to be interpreted with caution, due to the low degree of market openness to foreign investment.* Nevertheless, the TVECMs give similar and significant ECT coefficients in the unusual regime for both series, showing no space for arbitrage, hence no profitable investments.

* Financial indicators provided by the Institutional Profiles Database (CEPII 2009) show that, in the period considered, the Chinese market and, to a lesser extent, the Chilean market, exhibited a low level of financial openness compared with France, the United States and Australia.

9.6 Conclusion

This paper makes use of the Mediobanca Global Wine Industry Share Price Index with the aim of considering the wine market as a possible financial investment. The analyses include five wine-producing countries, Australia, Chile, China, France and the U.S., between January 1, 2001 to the end of February 2009.

A first-step investigation of the Mediobanca Index returns and abnormal returns shows that, apart from Australia, whose wine industry suffered a financial crisis during the mid-2000s, in all countries considered the wine indexes outperformed the composite indexes, revealing investment in wine stocks as a profitable investment *per se*.

We then focus our analyses on the long-run relationship between wine share price indexes and general stock market indexes in the same producing countries, examining their speed of adjustment to the long-run equilibrium. This is done using non-linear cointegration to capture price adjustments, which are activated when deviations from the equilibrium values exceed some threshold.

The results confirm the existence of threshold cointegration between wine and composite indexes for the period under study for all countries except Australia. In particular, in more mature markets (i.e. France and the US), when the gap between the wine index and the composite index exceeds a critical threshold, the speed of adjustment of the wine index is slower than that of the composite index. This means that wine price deviations from equilibrium last a longer time and informed traders can anticipate wine price movements over the short run and make profitable investments as a result of the weak-form efficiency and different speeds of adjustment.

In less mature markets (i.e. Chile and China), the wine index and the composite index are still non-linearly cointegrated, but there is no marked difference in the speed of adjustment between composite and wine prices. In those countries, results from threshold cointegration analysis lead to the conclusion that these markets are still not sufficiently effective to allow certain types of trading strategies.

Although the results may need to be interpreted with caution, the evidence from Chile and China is likely to be the consequence of the different economic situations that characterize the different countries in the analysis, which include a different level of development of financial markets, and a different level of market openness, as outlined earlier. Erb *et al.* (1997) and Garten (1997) have already noted that, in general, emerging markets are complex and a proper understanding of the dynamics they experience hinges upon many factors. In particular, in an emerging context, financial markets may be 'thin': their size may be comparatively small in terms of market capitalization, number of listed companies and trading volumes. Furthermore, they may be characterized by a different level of free-market capitalism and democracy, as well as other specific factors that shape the economic and political context that sets them apart within the countries under analysis.

References

- Anderson, H.M., Transaction costs and nonlinear adjustment toward equilibrium in the US Treasury Bill markets. *Oxford Bull. Econ. Statist.*, 1997, **59**, 465–484.
- Andrews, D.W.K., Testing for parameter instability and structural change with unknown change point. *Econometrica*, 1993, **61**, 821–856.
- Aslanidis, N. and Kouretas, G.P., Testing for two-regime threshold cointegration in the parallel and official markets for foreign currency in Greece. *Econ. Model.*, 2005, **22**(4), 665–682.
- Aylward, D.K., A documentary of innovation support among New World wine industries. *J. Wine Res.*, 2003, **14**(1), 31–43.
- Aylward, D.K. and Turpin, T., New wine in old bottles: A case study of innovation territories in 'New World' wine production. *Int. J. Innov. Mgmt.*, 2003, **7**(4), 501–525.
- Balke, N.S. and Fomby, T.B., Threshold cointegration. *Int. Econ. Rev.*, 1997, **38**, 627–646.

- Brooks, C., Rew, A.G. and Ritson, S., A trading strategy based on the lead–lag relationship between the spot index and futures contract for the FTSE 100. *Int. J. Forecast.*, 2001, **17**, 31–44.
- Bukenya, J. and Labys, W.C., Do fluctuations in wine stocks affect wine prices? Working Paper No. 37317, American Association of Wine Economists, 2007.
- Burton, J. and Jacobsen, J.P., The rate of return on wine investment. *Econ. Inquiry*, 2001, **39**, 337–350.
- Buyuksahin, B., Haigh, M.S. and Robe, M.A., Commodities and equities: Ever a “Market of one”? *J. Altern. Invest.*, 2010, **12**, 76–95.
- Campbell, G. and Guibert, N., Old World strategies against New World competition in a globalising wine industry. *Br. Food J.*, 2006, **108**(4), 233–242.
- Canning, P. and Perez, A., Economic geography of the U.S. wine industry. Working Paper, No. 22, American Association of Wine Economists, 2008.
- CEPII, Institutional Profiles Database, 2009. Available online at: <http://www.cepii.fr>
- Chong, J. and Miffre, J., Conditional correlation and volatility in commodity futures and traditional asset markets. *J. Altern. Invest.*, 2010, **12**, 61–75.
- Crain, S.J. and Lee, J.H., Volatility in wheat spot and futures markets, 1950–1993: Government farm programs, seasonality, and causality. *J. Finance*, 1996, **51**, 325–344.
- Dawson, P.J., Sanjuan, A.I. and White, B., Structural breaks and the relationship between barley and wheat futures prices on the London International Financial Futures Exchange. *Rev. Agric. Econ.*, 2010, **28**, 585–594.
- De Vittorio, A. and Ginsburgh, V., Pricing red wines of Medoc vintages from 1949 to 1989 at Christie’s auctions. *J. Soc. Statist. Paris*, 1996, **137**, 19–49.
- Dwyer, G. and Wallace, M., Cointegration and market efficiency. *J. Int. Money Finance*, 1992, **11**, 318–327.
- Enders, W. and Siklo, P.L., Cointegration and threshold adjustment. *J. Bus. Econ. Statist.*, 2001, **19**, 166–176.
- Erb, C.B. and Harvey, C.R., The strategic and tactical value of commodity futures. *Financ. Anal. J.*, 2006, **62**, 69–97.
- Erb, C.B., Harvey, C.R. and Viskanta, T.E., Demographics and international investments. *Financ. Anal. J.*, 1997, **53**(4), 14–28.
- Fama, E.F., The behaviour of stock market prices. *J. Bus.*, 1965, **38**, 31–105.
- Fogarty, J.J., The return to Australian fine wine. *Eur. Rev. Agric. Econ.*, 2006, **33**, 542–561.
- Fuertes, A.M., Miffre, J. and Rallis, G., Tactical allocation in commodity futures markets: Combining momentum and term structure signals. *J. Bank. Finance*, 2010, **34**, 2530–2548.
- Garten, J., Troubles ahead in emerging markets. *Harv. Bus. Rev.*, 1997, **75**, 38–50.
- German, H. and Kharoubi, C., WTI crude oil futures in portfolio diversification: The time-to-maturity effect. *J. Bank. Finance*, 2008, **32**, 2553–2559.
- Giuliani, E. and Bell, M., The micro-determinants of meso-level learning and innovation: Evidence from a Chilean wine cluster. *Res. Policy*, 2005, **34**(1), 47–68.
- Goodhue, R.E., Green, R.D., Heien, D.M. and Martin, P.L., California wine industry evolving to compete in 21st century. *Calif. Agric.*, 2008, **62**(1), 12–18.
- Gorton, G.B., Hayashi, F. and Rouwenhorst, G.K., The fundamentals of commodity futures returns. Working paper, National Bureau of Economic Research, 2007.
- Gorton, G.B. and Rouwenhorst, G.K., Facts and fantasies about commodity futures. *Financ. Anal. J.*, 2006, **62**, 47–68.
- Granger, C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969, **37**, 424–438.
- Granger, C.W.J., Some properties of time series data and their use in econometric model specification. *J. Econometr.*, 1981, **16**, 121–130.
- Gwynne, R., UK retail concentration, Chilean wine producers and value chains. *Geograph. J.*, 2008, **174**(2), 97–108.
- Hamilton, J.D., A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 1989, **57**(2), 357–384.

- Hansen, B.E., Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica*, 1996, **64**, 413–430.
- Hansen, B.E. and Seo, B., Testing for two-regime threshold cointegration in vector error-correction models. *J. Econometr.*, 2002, **110**, 293–318.
- Hokkio, C. and Rush, M., Market efficiency and cointegration: An application to Sterling and Deutschmark exchange rates. *J. Int. Money Finance*, 1989, **8**, 75–88.
- Hong, H. and Yogo, M., What does futures market interest tell us about the macroeconomy and asset prices? Working paper, Princeton University, 2011.
- Hussain, M., Cholette, S. and Castaldi, R., An analysis of globalization forces in the wine industry: Implications and recommendations for wineries. *J. Glob. Mktg.*, 2008, **21**(1), 33–47.
- Insel, B., The U.S. wine industry. *Business Econ.*, 2008, January.
- Jaeger, E., To save or savor: The rate of return to storing wine: Comment. *J. Polit. Econ.*, 1981, **89**, 584–592.
- Jawadi, F. and Koubaa, Y., Threshold cointegration between stock returns: an application of STECM Models. Working paper series, University of Paris X-Nanterre, 2004.
- Jenster, P. and Cheng, Y., Dragon wine: Developments in the Chinese wine industry. *Int. J. Wine Bus. Res.*, 2008, **20**(3), 244–259.
- Johansen, S., Statistical analysis of cointegrating vectors. *J. Econ. Dynam. Control*, 1988, **12**, 231–254.
- Johansen, S. and Juselius, K., Maximum likelihood estimation and inference on cointegration – with application to the demand for money. *Oxford Bull. Econ. Statist.*, 1990, **52**, 169–210.
- Jones, G.V. and Storchmann, K.H., Wine market prices and investment under uncertainty: An econometric model for Bordeaux Crus Classés. *Agric. Econ.*, 2001, **26**(2), 115–133.
- Jovanovic, B., Bubbles in prices of exhaustible resources. Working Paper No. 13320, National Bureau of Economic Research, Cambridge, MA, 2007.
- Krasker, W.S., The rate of return to storing wine. *J. Polit. Econ.*, 1979, **87**, 1363–1367.
- Krolzig, H.M., Markov-switching Vector Autoregressions. Modelling, Statistical Inference and Applications to Business Cycle Analysis, 1997 (Springer: Berlin).
- Lo, M. and Zivot, E., Threshold cointegration and nonlinear adjustment to the Law of One Price. *Macroecon. Dynam.*, 2001, **5**, 533–576.
- Maasoumi, E. and Racine, J., Entropy and predictability of stock market returns. *J. Econometr.*, 2002, **107**, 291–312.
- Masih, A. and Masih, R., Dynamic linkages and the propagation mechanism driving major international stock markets: An analysis of pre- and post-crash eras. *Q. Rev. Econ. Finance*, 1997, **37**, 859–885.
- Masset, P. and Henderson, C., Wine as an alternative asset class. Price evolution over the period 1996–2007, market segmentation and portfolio diversification. Paper presented at the Second Annual Conference, American Association of Wine Economists (AAWE), Portland, Oregon, 2008.
- Masset, P. and Weisskopf, J.P., Raise your glass: Wine investments and the financial crises. SSRN Research Papers, 2010.
- McMillan, D.G., Non-linear predictability of UK stock market returns. *Oxford Bull. Econ. Statist.*, 2003, **65**, 557–573.
- McMillan, D.G., Non-linear dynamics in international stock market returns. *Rev. Financ. Econ.*, 2005, **14**, 81–91.
- Miffre, J. and Rallis, G., Momentum strategies in commodity futures markets. *J. Bank. Finance*, 2007, **31**, 1863–1886.
- Mitry, D.J., Smith, D.E. and Jenster, P.V., China's role in global competition in the wine industry: A new contestant and future trends. *Int. J. Wine Res.*, 2009, **1**, 19–25.
- Patra, T. and Poshakwale, S.S., Long-run and short-run relationship between the main stock indexes: Evidence from the Athens stock exchange. *Appl. Financ. Econ.*, 2008, **18**, 1401–1410.
- Perez-Quiros, G. and Timmermann, A., Firm size and cyclical variations in stock returns. *J. Finance*, 2000, **55**, 1229–1262.

- Prat, G. and Jawadi, F., Nonlinear stock price adjustment in the G7 countries. Working paper, University of Paris Ouest La Defense Nanterre, 2009.
- Sanning, L.W., Shaffer, S. and Sharratt, J.M., Alternative investments: The case of wine. Working Paper No. 37322, American Association of Wine Economists, 2007.
- Schroeder, T.C. and Goodwin, B.K., Price discovery and cointegration for live hogs. *J. Fut. Mkts.*, 1991, **11**, 685–696.
- Schwarz, T.V. and Szakmary, A.C., Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. *J. Fut. Mkts.*, 1994, **14**, 147–167.
- Sercu, P., Uppal, R. and Van Hulle, C., The exchange rate in the presence of transaction costs: Implications for tests of purchasing power parity. *J. Finance*, 1995, **50**(4), 1309–1319.
- Shively, P.A., The nonlinear dynamics of stock prices. *Q. Rev. Econ. Finance*, 2003, **43**, 505–517.
- Shleifer, A., Inefficient Markets. An Introduction to Behavioural Finance, 2000 (Clarendon Lectures in Economics) (Oxford University Press: Oxford).
- Siourounis, G.D., Modelling volatility and test for efficiency in emerging capital markets: The case of the Athens stock exchange. *Appl. Financ. Econ.*, 2002, **1**, 47–55.
- Stoll, H.R. and Whaley, R.E., Commodity index investing and commodity futures prices. *J. Appl. Finance*, 2010, **20**(1), 7–46.
- Subramanian, U., Cointegration of stock markets in East Asia in the boom and early crisis period. Working Paper Series, Addis Ababa University, 2009.
- Szakmary, A.C., Shen, Q. and Sharma, S.C., Trend-following trading strategies in commodity futures: A re-examination. *J. Bank. Finance*, 2010, **34**, 409–426.
- Terblanche, N.S., Simon, E. and Taddei, J.-C., The need for a marketing reform: The wines of the Loire region. *J. Int. Food Agribus. Mktg*, 2008, **20**(4), 113–138.
- Tong, H. and Lim, K.S., Threshold autoregression, limit cycles and cyclical data. *J. R. Statist. Soc., Ser. B: Methodological*, 1980, **42**, 245–292.
- Tsay, R.S., Testing and modelling threshold autoregressive process. *J. Am. Statist. Assoc.*, 1989, **84**, 231–240.
- Tsay, R.S., Testing and modelling multivariate threshold models. *J. Am. Statist. Assoc.*, 1998, **93**, 1188–1202.
- Weil, R.L., Do not invest in wine, at least in the U.S. unless you plan to drink it, and maybe not even then. Paper presented at the 2nd International Conference of the Vineyard Data Quantification Society, Verona, Italy, 1993.
- Yang, J., Bessler, D.A. and Leatham, D.J., Asset storability and price discovery in commodity futures markets: A new look. *J. Fut. Mkts.*, 2001, **21**, 279–300.

10

Cross-Market Soybean Futures Price Discovery: Does the Dalian Commodity Exchange Affect the Chicago Board of Trade?

Liyan Han
Rong Liang
Ke Tang

10.1	Introduction	192
10.2	The Data	194
	The Construction of Continuous Price Series • Calculation of Futures Returns	
10.3	The SVAR Model.....	197
	Model Setup • Estimation Results • Impulse Response Analysis • Forecasted Error Variance Decomposition	
10.4	The VEC Model.....	205
	Model Setup • Estimation Results • Impulse Response and FEVD	
10.5	Conclusion	210
	Acknowledgements	210
	References.....	210

In this paper, we examine the role that the Dalian Commodity Exchange (DCE) plays in the global price discovery of soybean futures. We employ Structural Vector Autoregressive and Vector Error Correction models on the returns of the DCE and the Chicago Board of Trade (CBOT) soybean futures during trading and non-trading hours, and the result suggests that information transfers between DCE and CBOT in both directions, in particular that the DCE soybean futures prices influence price discovery in CBOT. Furthermore, the impulse response analysis and forecasted error variance decomposition show that the information transfer from DCE to CBOT is at a similar magnitude as that from CBOT to DCE. This shows that the DCE plays a prominent role in the global soybean futures price discovery. This conclusion differs from much of the literature, which mainly shows that the DCE is a satellite market and is dominated by the CBOT.

Keywords: Commodity markets; Commodity prices; Information in capital markets; Informational role of stock

JEL Classification: G1, G12, G13

10.1 Introduction

Soybeans, a source of high-quality vegetable protein, are widely used as food, fodder and industrial raw materials in the fields of health care and biodiesel. However, due to the limited amount of arable land and production technology, together with China's lack of transgenic varieties, China's soybean production has been restricted. However, China's demand for soybeans has been increasing, leading to an excess of soybean imports over its production since 2001. In 2008–2009, China's soybean consumption was 51.21 million tons, while its production was only 15.55 million tons; the remaining amount was imported. Figure 10.1 demonstrates trends in China's soybean production, consumption and import behaviour from 1994 to 2010.

Soybean futures, as one of the most actively traded species of agricultural futures, play a key role in the price discovery and hedging of soybeans and related products. Because soybean futures are simultaneously traded in many markets globally, the price formation in one market is not only affected by its domestic market information, but also by the related international market information. The information of one market can also be transferred to other markets. However, because different markets have different information processing capabilities, transaction costs, regulations and liquidity, even if they are information-linked, their relative contribution to the price discovery is not the same.

Currently, the world's major soybean futures markets include the Chicago Board of Trade (CBOT) of the United States, China's Dalian Commodity Exchange (DCE) and Japan's Tokyo Grain Exchange (TGE). The CBOT holds an important position in the soybean futures trading around the world. Over the past 10 years, the average trading volume of the CBOT soybean futures has been close to 20 million lots per year. Being highly liquid, the CBOT soybean market provides a convenient channel for investing and hedging, thus attracting not only the soybean-related production and circulation enterprises in and out of the United States, but also the international and domestic investment and speculative funds. Currently, there are two soybean futures contracts traded in the CBOT: one is a large-scale contract with a trading unit of 5000 bushels, and the other is a mini-contract with a trading unit of 1000 bushels. The DCE in China is another important soybean futures market in the world, at least in terms of trading




FIGURE 10.1 The time series of China's soybean production, imports and consumption. The sample period is 1994 to 2010.

volume. In 2008, the trading volume of soybean futures in the DCE was ranked second globally, surpassed only by that in the CBOT. The first trade of soybean futures contracts in the DCE can be traced back to 1993. The underlying product of that contract is soybeans, and the delivery goods can be both genetically modified soybeans and non-genetically modified soybeans. In 2001, the national agriculture policy and related management measures for genetically modified products were introduced, in line with which the DCE soybean contract was split into No. 1 yellow soybean contracts and No. 2 yellow soybean contracts in 2002. The No. 1 yellow soybean contract only allows the delivery of non-genetically modified soybeans and started trading on March 15, 2002. This contract represents the overall condition of China's domestic soybeans and is currently the world's largest-traded non-genetically modified soybean futures. The No. 2 yellow soybean contract started trading on December 22, 2004, and the delivery goods are genetically modified soybeans. However, its turnover is much smaller compared with the No. 1 contract, and there is thus a lack of liquidity. Table 10.1 provides the details of the No. 1 yellow soybean contracts of the DCE and the soybean futures of the CBOT, with which we can easily compare the two contracts.

China, as the major soybean import country in the world, may have played a key role in the global price discovery of soybean, which is the main issue that we investigate in this paper. By examining the relationship between the Chinese and the global soybean futures market, we can better understand China's position in the pricing of soybean futures, which can provide guidance for production, trade and policy making decisions.

Our paper focuses on the impact of the DCE on the CBOT in the price discovery process of soybean futures. We first explore the relationship between the DCE and the CBOT futures trading and non-trading hour returns with a Structural Vector Auto-Regression (SVAR) model, which can capture the lead-lag relationship between the two markets. The results show that the impacts of both the DCE on the CBOT and the CBOT on the DCE are significant and of a similar magnitude. A Vector Error Correction (VEC) model on the closing price of the soybean futures of the two markets further confirms this conclusion, which suggests that the closing prices of the DCE and the CBOT are cointegrated and

TABLE 10.1 Details of The DCE No.1 Soybean Futures Contract and The CBOT Soybean Futures Contract

	Dalian Commodity Exchange	Chicago Board of Trade
Product	No.1 soybeans	Soybeans
Trading unit	10MT/contract	5000 bushels (~136 metric tons)
Price quote	CNY/MT	Cents per bushel
Tick size	1 CNY/MT	1/4 of one cent per bushel
Daily price limit	4% of last settlement price (temporarily 5%)	\$0.70 per bushel expandable to \$1.05 and then to \$1.60 when the market closes at limit bid or limit offer
Contract months	Jan, Mar, May, July, Sep, Nov	Jan, Mar, May, July, Aug, Sep, Nov
Trading hours	9:00–11:30 a.m., 1:30–3:00p.m. Beijing Time, Monday–Friday	9:30 a.m.–1:15 p.m. Central Time, Monday–Friday
Last trading day	10th trading day of the delivery month	The business day prior to the 15th calendar day of the contract month
Last delivery day	7th day after the last trading day of the delivery month	Second business day following the last trading day of the delivery month
Deliverable grades	Quality standards of No.1 soybeans	#2 Yellow at contract price, #1 Yellow at a 6 cent/bushel premium, #3 Yellow at a 6 cent/bushel discount
Delivery location	The warehouses appointed by the DCE	The warehouses appointed by the CBOT

that both the CBOT and the DCE adjust according to this long-term relationship. These results signify that information flows from the DCE to the CBOT, and *vice versa*. This finding differs from the results of many existing studies, which mainly show that the CBOT is the dominant market and that the DCE is a satellite market. For example, previous research on this subject, such as Fung *et al.* (2003), find that the DCE is in a subordinate position in the price discovery process, i.e. the CBOT influences the DCE, but not *vice versa*. In this paper, using more *recent* data from 2002 to 2011, we find a different result, i.e. that the DCE significantly influences the CBOT in the soybean futures price discovery.

The research of cross-market price discovery was first proposed by Garbade and Silber (1979). Based on a study of the New York Stock Exchange and regional exchanges, the authors propose the concept of the 'dominant market' and the 'satellite market'. They posit that the regional exchanges are the satellite markets and that the price changes of the satellite markets lag behind those of the dominant markets. In other words, new information is first reflected in the price of the dominant market and, then, in the price of the satellite markets. Harris *et al.* (1995) examine the transactions of the IBM stock on the New York Stock Exchange, Pacific Stock Exchange and Midwest Stock Exchange with an error correction model, and the result shows that there is a long-term cointegration relationship among all three markets. Kleidon and Werner (1996) study the intraday patterns of volatility, volume and spreads of the stocks that are cross-listed in the UK and New York Stock Exchange, and they find that the order flow for the cross-listed securities are not fully integrated. Eun and Sabherwal (2003) study the Canadian stocks that are traded both on the Toronto Stock Exchange and the U.S. exchange. The authors find that prices are cointegrated in these two markets and that price adjustments exist in both markets. The role that the U.S. exchange plays in the price discovery of these shares is weaker overall than that in the Toronto Stock Exchange, but for some stocks, the U.S. exchange has a larger contribution. Grammig *et al.* (2005) investigate the price discovery of three stocks of German companies that are simultaneously listed on the New York Stock Exchange and Frankfurt Stock Exchange. They conclude that most price discovery first occurs in the Frankfurt market, and the New York Stock Exchange only plays a minor role in the price discovery of these stocks.

Regarding commodity futures, Booth and Ciner (1997) study the price series of the commodity futures from 1993 to 1995 in the Japan TGE and the U.S. CBOT. With a dynamic VAR model, they examine the relationship between the futures market prices of the two countries and the spillover effects of price volatility. The results show that the CBOT corn futures are dominant in the transmission of price information, but the TGE can also affect the prices of the CBOT commodity futures. Tse and Booth (1997) investigate the information transmission of the New York and London oil futures market with a Vector Error Correction (VEC) model and find that the New York oil futures market takes a dominant position in the information transmission process.

Our paper proceeds as follows. Section 10.2 presents the data. [Section 10.3](#) employs an SVAR model to test the information transfer between the DCE and CBOT soybean markets. As a robustness check, [Section 10.4](#) utilizes a VEC model to investigate the price adjustment among two exchanges. [Section 10.5](#) concludes the paper.

10.2 The Data

The selected data of the DCE No. 1 soybean futures prices and of the CBOT soybean futures prices are dated from March 2002, when the No. 1 yellow soybeans were first launched in China, to September 2011. We choose the No. 1 soybean futures from the DCE because the No. 1 yellow soybean contract has a much larger trading volume and, hence, more liquidity than the No. 2 contract.

10.2.1 The Construction of Continuous Price Series

Because futures contracts of different months are traded simultaneously, it is of practical importance to choose the one that best represents the commodity prices. In previous studies, there are generally

two approaches. One is to select the price of the front-month futures contract and, when the front-month contract enters the delivery month, to select the next front-month contract as the representative. Another approach focuses on the trading volume and selects the contract with the largest trading volume as the representative (Booth and Ciner 1997).

For the CBOT soybean futures, the front-month contract is generally the most actively traded; therefore, prices (returns) of the front-month contract are used for the CBOT. In China, however, the front-month soybean futures contract is not the most actively traded. It is the contracts of January, May and September that are generally more actively traded due to the seasonal soybean production and circulation pattern. Nonetheless, the second approach means that the price series may switch between the front contract, the second front contract and possibly the third front contract, which could result in a mismatch issue with the CBOT price series. Therefore, we still draw on the first approach for the DCE and select the prices (returns) of the front-month contract. We obtain the daily DCE and CBOT futures prices from the Wind database and Datastream, respectively.

Because holidays in the United States and China are not exactly the same and futures trading stops during the holidays, the trading dates of soybean futures in the DCE and those in the CBOT do not overlap entirely. Following Tse *et al.* (1996) and Xu and Fung (2005), we remove the days for which transactions only exist in one market and construct the price (return) series with a complete match between these two markets. The prices of the DCE futures are adjusted to U.S. dollars using the daily exchange rate from the Wind database. Figure 10.2 depicts the sequence of the normalized closing prices of the DCE soybean futures contract and the CBOT soybean futures contract. Note that the DCE price series is converted into US dollar denomination with the exchange rate from the Wind database. The two price series show a similar trend.

10.2.2 Calculation of Futures Returns

Before we present the calculation of different futures returns, it should be noted that the trading times of the two markets do *not* overlap. Due to the time difference, the calendar time of the Chinese market is ahead of the U.S. market, and when the DCE is trading, the CBOT is closed, and *vice versa*. Specifically, under the standard time, there is a 14-hour time difference between the DCE and the CBOT, while




FIGURE 10.2 The time series of the price of the DCE No.1 soybean futures and the CBOT soybean futures. The price series of DCE soybean futures is converted into US dollars with the corresponding exchange rate. The futures prices are normalized to be 100 with March 15, 2002 as the base period. The sample starts on March 15, 2002 and ends on September 2, 2011.

under the daylight saving time, the time difference is 13 hours. Under the standard time, the DCE opens at 9:00 Beijing time (Chicago time previous day 19:00) and closes at 15:00 Beijing time (Chicago time 1:00); the CBOT opens at 23:30 Beijing time (Chicago time 9:30) and closes at 3:10 Beijing time (Chicago time 13:10). During the daylight saving time period, the case is similar and there is still no overlap of the trading time. Figure 10.3 shows the time line of the non-trading and trading hours for both the CBOT and DCE markets. To examine the price response to information, two classes of returns should be considered: trading hour returns and non-trading hour returns. For example, the information generated in the trading hours of the Chinese market may affect the non-trading hour returns of the U.S. market, i.e. the U.S. market's opening price at date t may reflect digested information generated from the Chinese market on date t .

The trading and non-trading hour returns at date t for the DCE and the CBOT are defined as follows:

DCE soybean trading hour returns: $CN_INR_t = CN_CLOSE_t - CN_OPEN_t$,

DCE soybean non-trading hour returns: $CN_OUTR_t = CN_OPEN_t - CN_CLOSE_{t-1}$,

CBOT soybean trading hour returns: $US_INR_t = US_CLOSE_t - US_OPEN_t$,

CBOT soybean non-trading hour returns: $US_OUTR_t = US_OPEN_t - US_CLOSE_{t-1}$,

where CN_INR_t is the trading hour returns of the DCE soybean futures at date t , CN_OUTR_t is the non-trading hour returns of the DCE soybean futures at date t , CN_CLOSE_t is the log of the closing price of the DCE soybean futures at date t , and CN_OPEN_t is the log of the opening price of the DCE soybean futures at date t . Similarly, the trading and non-trading hour returns for the CBOT soybean futures are US_INR_t and US_OUTR_t , and US_CLOSE_t and US_OPEN_t correspond to the log of the closing price and opening price, respectively, of the CBOT soybean futures at date t . Panel A of Table 10.2 gives the respective summary statistics of the different returns for DCE and CBOT soybean futures.

From panel A, we can see that the DCE and the CBOT returns, on average, are zero, negatively skewed and relatively thin-tailed. Before estimating the SVAR model, we need to first test the stationarity of the return series. To ensure a sound conclusion, ADF, PP and KPSS tests are used for the unit root test. Note that we exclude both trend and intercept in the test since the plot (not given here) indicates that the return series seem to fluctuate around zero. Panel B of Table 10.2 gives the test result. All three tests indicate that the four return series are all stationary. In this case, we need not consider the possible cointegration problem in the Structural VAR model.




FIGURE 10.3 The time difference of the DCE and the CBOT exchanges. It shows the relative length of the trading time and non-trading time, as well as the correspondence of the DCE and the CBOT trading and non-trading time due to the time difference.

TABLE 10.2 Summary Statistics and the Unit Root Test of the Return of the DCE No.1 Soybean Futures and the Cbot Soybean Futures. Panel A Gives the Summary Statistics, Where the First Two Columns Present the Statistics of the Trading and Non-Trading Hour Returns of the DCE No. 1 Soybean Futures While the Last Two Columns Give Statistics of the Trading and Non-Trading Hour Returns of the CBOT Futures. Panel B Presents the Unit Root Test Result of the ADF, PP And KPSS Tests. Values in Parentheses are the Corresponding *P*-Values. The Sample Period is March 15, 2002 to September 2, 2011

Test statistic	CN_INR	CN_OUTR	US_INR	US_OUTR
Panel A: Summary statistics				
Mean	0.00024 (0.2581)	0.00021 (0.399)	6.08E-05 (0.8234)	0.00042 (0.1195)
Median	0	4.46E-05	0.00039	0.00050
Max	0.08744	0.08662	0.07633	0.0734
Min	-0.07599	-0.12867	-0.074	-0.23307
Std. dev.	0.0104	0.0119	0.0131	0.0131
Skewness	-0.115	-0.307	-0.285	-2.586
Kurtosis	12.53	15.49	5.21	49.88
Jarque-Bera	8819.6 (0.000)	15,178.7 (0.000)	504.2 (0.000)	215,954 (0.000)
Observations	2330	2330	2330	2330
Panel B: Unit root test				
ADF	-47.71 (0.0001)	-48.79 (0.0001)	-49.11 (0.0001)	-48.93 (0.0001)
PP	-4.72 (0.0001)	-48.80 (0.0001)	-49.16 (0.0001)	-48.98 (0.0001)
KPSS	0.245	0.124	0.146	0.097

Note: For the KPSS test, all the series are tested including the intercept; critical values at the 10%, 5%, and 1% significance levels are 0.347, 0.463, and 0.7439, respectively.

10.3 The SVAR Model

This section investigates the relationship between the DCE and the CBOT in soybean futures price discovery. In previous research on cross-market price discovery with non-overlapping data, the primary method has been the VAR model with some restrictions on the coefficients or the VEC model. However, the SVAR model can better capture the contemporaneous relationship among the related variables. Hence, we first use the SVAR model to explore the relationship between DCE and CBOT soybean futures. We now report the model setup and estimation result.

10.3.1 Model Setup

The SVAR model can capture the interaction between variables of the same period. A SVAR model can be expressed as*

$$\mathbf{B}\mathbf{y}_t = \Gamma_0 + \Gamma_1\mathbf{y}_{t-1} + \Gamma_2\mathbf{y}_{t-2} + \dots + \Gamma_p\mathbf{y}_{t-p} + \boldsymbol{\mu}_t, \quad t = 1, 2, \dots, T, \quad (10.1)$$

where matrix $\mathbf{B}$ captures the relationship between variables of the same period. $\boldsymbol{\mu}_t$ is i.i.d. with a mean of zero and is known as the structural disturbance vector. If the matrix $\mathbf{B}$ is invertible, the model can be simplified as follows:

$$\mathbf{y}_t = \mathbf{B}^{-1}\Gamma_0 + \mathbf{B}^{-1}\Gamma_1\mathbf{y}_{t-1} + \mathbf{B}^{-1}\Gamma_2\mathbf{y}_{t-2} + \dots + \mathbf{B}^{-1}\Gamma_p\mathbf{y}_{t-p} + \mathbf{B}^{-1}\boldsymbol{\mu}_t, t = 1, 2, \dots, T. \quad (10.2)$$

* Note that bold letters denote a matrix or vector in this paper.

After imposing some constraints on the matrix $\mathbf{B}$, we can estimate the SVAR model and obtain the identifiable structural impact.

For the SVAR model of the four return series of the DCE soybean futures and CBOT soybean futures, we define $\mathbf{y}_t$ as follows:

$$\mathbf{y}_t := \begin{pmatrix} CN_OUTR_t \\ CN_INR_t \\ US_OUTR_t \\ US_INR_t \end{pmatrix}.$$

Since this SVAR model contains four vectors, we need six constraints to identify it. When making constraints on matrix $\mathbf{B}$, we follow the rule that returns appearing earlier can influence the returns appearing later. As shown in Figure 10.3, due to the time difference, the trading time of the DCE at date t corresponds to the non-trading time of the CBOT at date t . Because the duration of the non-trading time is longer than the duration of the trading time, the non-trading hour's returns of one market can be affected by the other market's trading hour returns of the same date, while the trading hour returns of one market are not subject to the influence of the other market's non-trading hour returns of the same date. For example, as illustrated in Figure 10.3, when the DCE closes at time t , the CBOT is not yet open; therefore, the opening price of the CBOT at day t can be affected by the DCE's trading time information at day t . Because the CBOT's non-trading hour returns are calculated as the CBOT's open price at time t divided by the previous day's closing price, it therefore may also be affected by the DCE's trading hour returns at day t . The relationships between the returns of other pairs of the DCE and the CBOT returns can be analysed in a similar way. With Equation (10.1) and the definition of vector $\mathbf{y}_t$, we thus impose the following restrictions according to the order of occurrence for the DCE and CBOT trading hour and non-trading hour returns.

First, neither the CBOT trading and non-trading hour returns nor the DCE trading hour returns can affect the DCE non-trading hour returns of the same day. Thus, $b_{12} = b_{13} = b_{14} = 0$.

Second, only DCE non-trading hour returns of the same date can affect the DCE trading hour returns. Note that the CBOT trading hour and non-trading hour returns have no effect on the same day's DCE trading hour returns because these two returns are not even known when the DCE closes at date t . Thus, $b_{23} = b_{24} = 0$.

Third, the DCE trading and non-trading hour returns of the same date can affect the CBOT non-trading hour return, while the CBOT trading hour returns of the same date have no effect on the CBOT non-trading hour returns. Thus, $b_{34} = 0$.

Fourth, the DCE trading and non-trading hour returns and the CBOT non-trading hour returns of the same date can all affect the CBOT trading hour returns.

In summary, $\mathbf{B}$ in the SVAR model can be expressed as

$$\mathbf{B} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -b_{21} & 1 & 0 & 0 \\ -b_{31} & -b_{32} & 1 & 0 \\ -b_{41} & -b_{42} & -b_{43} & 1 \end{pmatrix}.$$

Note that there is a minus sign in front of the coefficients. This is for the convenience of interpretation, which will soon become clear with the following explanation. After estimating the matrix $\mathbf{B}$, we can write the model in the structural form and obtain a sense of the impact of the DCE on the CBOT.

To make the interpretation easier, we can rearrange Equation (10.1) into the following form:

$$\mathbf{y}_t = \Gamma_0 + (\mathbf{I} - \mathbf{B})\mathbf{y}_t + \Gamma_1\mathbf{y}_{t-1} + \Gamma_2\mathbf{y}_{t-2} + \dots + \Gamma_p\mathbf{y}_{t-p} + \mu_t, \quad t = 1, 2, \dots, T \quad (10.3)$$

where

$$\mathbf{I} - \mathbf{B} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ b_{21} & 0 & 0 & 0 \\ b_{31} & b_{32} & 0 & 0 \\ b_{41} & b_{42} & b_{43} & 0 \end{pmatrix},$$

which shows the contemporaneous relationship of $\mathbf{y}_t$. This also explains why we set a minus sign in front of the coefficients of matrix $\mathbf{B}$. Collecting the relevant terms, we can write the model more explicitly as follows:

$$\begin{aligned} CN_OUTR_t &= a_1 + \sum_{i=1}^n c_{1i} CN_OUTR_{t-i} \\ &\quad + \sum_{i=1}^n d_{1i} CN_INR_{t-i} \\ &\quad + \sum_{i=1}^n f_{1i} US_OUTR_{t-i} \\ &\quad + \sum_{i=1}^n g_{1i} US_INR_{t-i} + h_1 Mon_t + \varepsilon_{1t}, \end{aligned} \quad (10.4)$$

$$\begin{aligned} CN_INR_t &= a_2 + \sum_{i=0}^n c_{2i} CN_OUTR_{t-i} \\ &\quad + \sum_{i=1}^n d_{2i} CN_INR_{t-i} + \sum_{i=1}^n f_{2i} US_OUTR_{t-i} \\ &\quad + \sum_{i=1}^n g_{2i} US_INR_{t-i} + h_2 Mon_t + \varepsilon_{2t}, \end{aligned} \quad (10.5)$$

$$\begin{aligned} US_OUTR_t &= a_3 + \sum_{i=0}^n c_{3i} CN_OUTR_{t-i} \\ &\quad + \sum_{i=0}^n d_{3i} CN_INR_{t-i} \\ &\quad + \sum_{i=1}^n f_{3i} US_OUTR_{t-i} \\ &\quad + \sum_{i=1}^n g_{3i} US_INR_{t-i} + h_3 Mon_t + \varepsilon_{3t}, \end{aligned} \quad (10.6)$$

$$\begin{aligned}
US_INR_t = & a_4 + \sum_{i=0}^n c_{4i} CN_OUTR_{t-i} \\
& + \sum_{i=0}^n d_{4i} CN_INR_{t-i} + \sum_{i=0}^n f_{4i} US_OUTR_{t-i} \\
& + \sum_{i=1}^n g_{4i} US_INR_{t-i} + h_4 Mon_t + \varepsilon_{4t}.
\end{aligned} \tag{10.7}$$

In the above equations, when the subscript i is zero, the contemporaneous value at date t of that explanatory variable can affect the left-hand side variable. For example, in (10.5), the summation symbol of CN_OUTR_t has a subscript of 0, indicating that the non-trading hour returns of the DCE soybean futures can affect its trading hour returns at the same date. The 0 subscripts in the other equations are interpreted similarly. Actually, it is easy to see that $c_{20} = b_{21}$, $c_{30} = b_{31}$, $d_{30} = b_{32}$, $c_{40} = b_{41}$, = and $f_{40} = b_{43}$. Also note that we include dummy variable Mon in the SVAR model, which equals 1 when the day is Monday, and 0 otherwise. This is to capture the well-documented weekend effect, that is, Monday returns tend to be significantly negative.

In addition, as [Figure 10.3](#) shows, although the CBOT market may not affect the DCE market of the same calendar date, the lag one period effect may still be significant. In other words, the information generated in the CBOT at date $t-1$ may affect the DCE at date t . This potential impact can be verified by examining the coefficients of Γ_1 (see [Equation \(10.3\)](#)).* For this model, Γ_1 is in the following form:

$$\Gamma_1 = \begin{pmatrix} \Gamma_{11} & \Gamma_{12} & \Gamma_{13} & \Gamma_{14} \\ \Gamma_{21} & \Gamma_{22} & \Gamma_{23} & \Gamma_{24} \\ \Gamma_{31} & \Gamma_{32} & \Gamma_{33} & \Gamma_{34} \\ \Gamma_{41} & \Gamma_{42} & \Gamma_{43} & \Gamma_{44} \end{pmatrix},$$

where Γ_{11} represents the influence of CN_OUTR_{t-1} on CN_OUTR_t , and Γ_{12} represents the influence of CN_INR_{t-1} on CN_OUTR_t . The other parameters can be interpreted similarly. The coefficients that we focus on are Γ_{13} , Γ_{14} , Γ_{23} and Γ_{24} . Here, Γ_{13} and Γ_{14} represent the effect of the $t-1$ non-trading and trading hour returns of the CBOT on the date t non-trading hour returns of the DCE, while Γ_{23} and Γ_{24} represent the effect of the $t-1$ non-trading and trading hour returns of the CBOT on the date t trading hour returns of the DCE. With a detailed analysis of Γ_1 and $\mathbf{B}$, we can better assess the information transfer between CBOT and DCE. Also note that Γ_{ij} actually equals c_{ij} in Equations (10.4)–(10.7).

10.3.2 Estimation Results

The BIC statistics indicate that the appropriate lag order of the SVAR model is 2. The detailed estimation results for the b_{ij} values are shown in [Table 10.3](#). As seen from the table, except for the fact that b_{41} and b_{42} are not significant, all of the coefficients are significant at the 1% significance level. Because we are interested in the impact of the two exchanges, here we focus on the analyses of b_{31} , b_{32} , b_{41} and b_{42} . Among these coefficients, b_{31} and b_{32} represent the impact of the non-trading and trading hour returns of the DCE soybean futures on the non-trading hour returns of the CBOT soybean futures, and b_{41} and b_{42} represent the impact of the DCE soybean futures on the trading hour returns of the CBOT soybean futures. It can be seen that both the non-trading and trading hour returns shock of the DCE soybean futures has a significant *positive* impact on the CBOT soybean futures non-trading hour returns, while

* In the empirical analysis, we allow a lag of order 2 in the SVAR; hence, Γ_2 also influences the DCE returns. However, because its magnitude is quite small from our empirical results, we omit it when presenting the value of Γ_2 .

TABLE 10.3 Estimated Results of Matrix **B**, Which Represents the Relationship between Variables of the Same Period in the SVAR Model. The Variables in the SVAR Model are in the order of *CN_OUTR*, *CN_INR*, *US_OUTR* and *US_INR*. The Values in Brackets are the Corresponding *T*-Statistics. *** and ** Denote the 1% and 5% Significance Level, Respectively. The Sample Period is March 15, 2002 to September 2, 2011

	Coefficient
b_{21}	-0.358*** [-19.88]
b_{31}	0.292*** [11.35]
b_{41}	0.0004 [0.015]
b_{32}	0.245*** [8.95]
b_{42}	0.034 [1.18]
b_{43}	-0.068*** [-3.21]

TABLE 10.4 Estimated Results of Matrix Γ_1 , Which represents The Lag One Period Effect of the Independent Variables on the Dependent Variables in the SVAR Model. The Variables in the SVAR Model are in the order of *CN_OUTR*, *CN_INR*, *US_OUTR* and *US_INR*. The Values in Brackets are the Corresponding *T*-Statistics. *** and ** Denote the 1% and 5% Significance Level, Respectively. The Sample Period is March 15, 2002 to September 2, 2011

Dependent variable	Independent variable			
	<i>CN_OUTR</i> _{<i>t-1</i>}	<i>CN_INR</i> _{<i>t-1</i>}	<i>US_OUTR</i> _{<i>t-1</i>}	<i>US_INR</i> _{<i>t-1</i>}
<i>CN_OUTR</i>	-0.086*** [-3.73]	-0.173*** [-7.16]	0.179*** [9.89]	0.296*** [17.02]
<i>CN_INR</i> _{<i>t</i>}	-0.028 [0.11]	-0.042 [0.86]	0.025** [-2.29]	0.080 [-1.58]
<i>US_OUTR</i> _{<i>t</i>}	0.045 [0.75]	0.027 [-0.65]	-0.059 [-0.77]	-0.032** [2.35]
<i>US_INR</i> _{<i>t</i>}	0.012 [0.39]	0.016 [0.61]	-0.027 [-1.27]	-0.017 [-1.01]

neither the non-trading nor trading hour returns of the DCE soybean futures have a significant influence on the CBOT soybean futures' trading hour returns. The reason why this happens may be that the time lag between the DCE returns and the CBOT soybean futures' trading hour return is relatively long, and the information of the DCE may have already been reflected in the opening price of the CBOT soybean futures.

Table 10.4 gives the estimated results for Γ_1 . As we can see, except for Γ_{24} , the estimated results for Γ_{13} , Γ_{14} and Γ_{23} are all significant at the 5% significance level. Both the non-trading and trading hour returns of the lag one date CBOT soybean futures have a significant positive impact on the DCE soybean futures' non-trading hour returns, and the non-trading hour returns of the lag one date CBOT soybean futures also have a significant positive impact on the DCE soybean futures' trading hour returns. This indicates that the CBOT prices influence the DCE prices, which is consistent with Fung *et al.* (2003).

To conclude, the DCE and the CBOT influence each other in both directions. It needs to be emphasized that the DCE has a significant effect on the CBOT price discovery. With China's large-scale and still increasing import demand, this result is reasonable.

10.3.3 Impulse Response Analysis

In the SVAR model, by orthogonalizing the impulse response function, we can investigate the impact of each variable on the other variables separately. The responses of the CBOT non-trading and trading hour returns to one shock of the DCE trading and non-trading hour returns are illustrated in Figure 10.4. It can be seen that the response of the CBOT non-trading hour returns to one unit of positive shock of the DCE non-trading hour returns in the first period is positive, but the response quickly dampens down in the second period. The pattern and the magnitude of the response to one unit of positive shock of the DCE trading hour returns are quite similar.

As shown in Figure 10.4, the response of the CBOT trading hour returns to one unit of positive shock of the DCE non-trading hour returns in the first period is only marginal, and soon dampens down in the following period. It is worth noting that the magnitude of the response here is much smaller than the response of the CBOT non-trading hour returns to the DCE shock. However, this is anticipated if we consider the time difference between the DCE returns and the CBOT trading hour returns.

To provide a more comprehensive picture, we now examine how the DCE responds to the CBOT non-trading and trading hour returns, which are illustrated in Figure 10.5. Consistent with our estimates of the coefficients of matrix Γ_p , the impulse response analysis in the upper panels of Figure 10.5 also indicates that the CBOT non-trading and trading hour returns shock will positively affect the DCE non-trading hour returns in the next period. The magnitude of the response of the DCE non-trading hour returns to the CBOT shock is similar to, albeit a little larger than, that of the CBOT non-trading hour returns to the DCE shock shown in Figure 10.4.




FIGURE 10.4 The responses of the CBOT non-trading and trading hour returns to one unit of positive shock of the DCE non-trading and trading hour returns. The sample period is March 15, 2002 to September 2, 2011.




FIGURE 10.5 The responses of the DCE non-trading and trading hour returns of the soybean futures to one unit of positive shock of the CBOT non-trading and trading hour returns. The sample period is March 15, 2002 to September 2, 2011.

The lower panels of Figure 10.5 show that the DCE trading hour returns respond negatively to the CBOT shock in the previous period. This finding is consistent with the estimation result of Γ_1 in Table 10.4. A reasonable explanation is that the opening price of the DCE soybean futures overreacts to the information from the CBOT, and the closing price of the DCE soybean futures reverts to the normal level, leading to a negative response of the DCE trading hour returns to the CBOT return shock.

To sum up, the impulse response analysis indicates that the impact of the DCE market on the CBOT is significantly positive, pointing to the reality that the China soybean futures market is playing an important role in global soybean price discovery.

10.3.4 Forecasted Error Variance Decomposition

In line with the cross-market price discovery study of Grammig *et al.* (2005), we also explore the ‘information share’ of each variable in the price discovery process, which is defined as the proportion of the innovation variance in the variable i that can be attributed to innovations in the variable j . The information shares are obtained by forecasted error variance decomposition (FEVD).

When there is contemporaneous correlation among the innovations, identifying an independent information share for each variable without further restrictions is not possible. However, the upper and lower bounds can be provided using the Cholesky factorization for the covariance matrix of the error term. The Cholesky factorization will give the upper bound of the share for the variable ordered first and a lower bound of the share for the variable ordered last.

However, since there are four variables in the SVAR model, it is impossible to give the upper bound and the lower bound of each variable at the same time while maintaining that the total information shares sum to 100. Thus, we consider two sets of ordering; one is as the order of the SVAR model, $D_OUTR\ D_INR\ US_OUTR\ US_INR$, which approximates the upper bound for the DCE market, and the other ordering is $US_OUTR\ US_INR\ D_OUTR\ D_INR$, which approximates the upper bound for the CBOT market.

Table 10.5 presents the results of the FEVD for the DCE non-trading and trading hour returns. As shown in the table, the non-trading and trading hour returns of the CBOT together contribute approximately 13–16% to the forecasting of the DCE non-trading hour returns. Their contribution to the DCE trading hour returns is much weaker, which is consistent with the result of the impulse response analysis. It is also worth noting that the fraction of the forecasting error of the DCE non-trading hour returns attributable to the returns of the CBOT significantly differs from zero in the second period, which is consistent with the time difference between the two markets. The variance decomposition for the CBOT non-trading and trading hour returns is shown in Table 10.6. Approximately 0.2–6% of

TABLE 10.5 Variance Decomposition of the Non-Trading and Trading Hour Return of the DCE No. 1 Soybean Futures. Panel A is the Variance Decomposition for the DCE Non-Trading Hour Return, While Panel B is for the DCE Trading Hour Return. In Each Panel, The First Column is the Period; The Second to the Fifth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to CN_OUTR , CN_INR , US_OUTR and US_INR , Respectively, with the Cholesky ordering of $CN_OUTR\ CN_INR\ US_OUTR\ US_INR$; The Sixth to the Ninth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to CN_OUTR , CN_INR , US_OUTR and US_INR , Respectively, with the Cholesky ordering of $US_OUTR\ US_INR\ CN_OUTR\ CN_INR$. The Sample Period is March 15, 2002 to September 2, 2011

Period	CN_OUTR	CN_INR	US_OUTR	US_INR	CN_OUTR	CN_INR	US_OUTR	US_INR
Panel A: Variance decomposition of CN_OUTR								
1	100	0	0	0	97.02	0.00	2.97	0.01
2	85.52	1	2.84	10.64	83.00	1.87	4.67	10.45
3	85.07	1.5	2.84	10.59	82.56	2.39	4.65	10.40
4	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
5	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
6	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
7	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
8	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
9	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
10	85.05	1.51	2.85	10.59	82.54	2.40	4.66	10.40
Panel B: Variance decomposition of CN_INR								
1	14.5	85.5	0	0	16.32	82.61	1.01	0.07
2	14.47	85.22	0.2	0.11	16.26	82.36	1.20	0.17
3	14.52	85.05	0.21	0.22	16.30	82.19	1.22	0.29
4	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
5	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
6	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
7	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
8	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
9	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
10	14.52	85.04	0.21	0.23	16.30	82.18	1.22	0.30
Cholesky ordering:					Cholesky ordering:			
$CN_OUTR\ CN_INR\ US_OUTR\ US_INR$					$US_OUTR\ US_INR\ CN_OUTR\ CN_INR$			

TABLE 10.6 Variance Decomposition of the Non-Trading and Trading Hour Return of the CBOT Soybean Futures. Panel A is the Variance Decomposition for the CBOT Non-Trading Hour Return, While Panel B is for the CBOT Trading Hour Return. In Each Panel, The First Column is the Period; The Second to the Fifth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_OUTR*, *CN_INR*, *US_OUTR* and *US_INR*, Respectively, with the Cholesky Ordering of *CN_OUTR* *CN_INR* *US_OUTR* *US_INR*; The Sixth to the Ninth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_OUTR*, *CN_INR*, *US_OUTR* and *US_INR*, Respectively, with the Cholesky ordering of *US_OUTR* *US_INR* *CN_OUTR* *CN_INR*. The Sample Period is March 15, 2002 to September 2, 2011

Period	<i>CN_OUTR</i>	<i>CN_INR</i>	<i>US_OUTR</i>	<i>US_INR</i>	<i>CN_OUTR</i>	<i>CN_INR</i>	<i>US_OUTR</i>	<i>US_INR</i>
Panel A: Variance decomposition of <i>US_OUTR</i>								
1	2.97	3.23	93.80	0.00	0.00	0.00	100.00	0.00
2	3.00	3.24	93.53	0.24	0.05	0.02	99.70	0.23
3	3.03	3.32	92.96	0.69	0.10	0.09	99.12	0.69
4	3.03	3.32	92.96	0.69	0.10	0.09	99.12	0.69
5	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
6	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
7	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
8	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
9	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
10	3.03	3.32	92.95	0.69	0.10	0.09	99.11	0.69
Panel B: Variance decomposition of <i>US_INR</i>								
1	0.05	0.02	0.44	99.50	0	0	0.43	99.57
2	0.05	0.02	0.50	99.43	0.001	0.016	0.48	99.50
3	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
4	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
5	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
6	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
7	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
8	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
9	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
10	0.06	0.06	0.60	99.27	0.003	0.034	0.62	99.35
Cholesky ordering:				Cholesky ordering:				
<i>CN_OUTR</i> <i>CN_INR</i> <i>US_OUTR</i> <i>US_INR</i>				<i>US_OUTR</i> <i>US_INR</i> <i>CN_OUTR</i> <i>CN_INR</i>				

the CBOT non-trading hour returns forecast error is attributable to the DCE trading and non-trading hour returns.

Combining the results from [Tables 10.5](#) and 10.6, the DCE and the CBOT both have a significant effect on each other, but the magnitude of the impact of the CBOT on the DCE is larger. It seems that the information transmission in the two markets is bidirectional, but the CBOT is still in the dominant position.

In the following section, as a robustness check, we use the VEC model to test the price interdependence between these two exchanges.

10.4 The VEC Model

In this section, we first present the model setup, we then show the estimation results, and, finally, we present an analysis of the results. Note that the VEC model emphasizes the long-term cointegration relationship, and hence we omit the short-term time difference between the two markets in the VEC model.

10.4.1 Model Setup

A cointegrated VAR system can be written in the following form:

$$\Delta \mathbf{y}_t = \zeta_1 \Delta \mathbf{y}_{t-1} + \zeta_2 \Delta \mathbf{y}_{t-2} + \dots + \zeta_{p-1} \Delta \mathbf{y}_{t-p+1} + \boldsymbol{\alpha} + \mathbf{K} \mathbf{A}' \mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \quad (10.8)$$

where $\mathbf{A}' \mathbf{y}_{t-1}$ is the long-term cointegration relationship between the series, while $\mathbf{K}$ reflects the short-term reaction to the deviation from this long-term relationship. The economic implication of a VEC model is that if there is a long-term equilibrium relationship between the relevant variables, the short-term variation of these variables can be viewed as a partial adjustment to the long-term relationship.

We employ a VEC model for the closing log price of the CBOT and the DCE to explore the information transmission mechanism between the two markets. The VEC model for the closing price of the CBOT and the DCE is as follows:

$$\begin{aligned} \Delta CN_CLOSE_t &= \sum_{i=1}^p \zeta_{1i} \Delta CN_CLOSE_{t-i} \\ &+ \sum_{i=1}^q \gamma_{1i} \Delta US_CLOSE_{t-i} + \alpha_1 \\ &+ k_1 (CN_CLOSE_{t-1} + a \\ &\times US_CLOSE_{t-1}) + h_1 \text{Mon}_t + \varepsilon_{1t}, \end{aligned} \quad (10.9)$$

$$\begin{aligned} \Delta US_CLOSE_t &= \sum_{i=1}^p \zeta_{2i} \Delta CN_CLOSE_{t-i} \\ &+ \sum_{i=1}^q \gamma_{2i} \Delta US_CLOSE_{t-i} + \alpha_2 \\ &+ k_2 (CN_CLOSE_{t-1} + a \\ &\times US_CLOSE_{t-1}) + h_2 \text{Mon}_t + \varepsilon_{2t}, \end{aligned} \quad (10.10)$$

where CN_CLOSE_t denotes the log of the closing price of the DCE soybean futures at day t , which is adjusted to U.S. dollars with the daily exchange rate from the Wind database, and US_CLOSE_t denotes the log of the closing price of the CBOT soybean futures at day t . $\Delta CN_CLOSE_t = CN_CLOSE_t - CN_CLOSE_{t-1}$, and the other first differences of the price series are defined similarly. In this model, $CN_CLOSE_t + a \times US_CLOSE_t$ reflects the long-term relationship between the closing prices of the two markets, and $k_i (i = 1, 2)$ reflects how each market adjusts to the long-term relationship. With this model, we can further test the relationship between the two markets. If k_1 is significant while k_2 is not, then the DCE reacts to the long-term relationship between the two markets while the CBOT does not, which means that the information is transmitted from CBOT to the DCE. Other cases can be interpreted similarly. Note that, as in the SVAR model, we also include dummy variable Mon to account for the widely discussed weekend effect.

10.4.2 Estimation Results

We first conduct the unit root test for the price series with three tests, the ADF, PP and KPSS tests. In this way, we can ensure a sounder conclusion since the ADF and PP tests have low test power. Since

the variables are in logarithms, a time trend, which implies an ever-increasing (or decreasing) rate of change in the first difference of the variable, is unlikely. Thus, we exclude the time trend but include an intercept in the unit root test. Panel A of Table 10.7 gives the test result. For both the series, the ADF and PP tests cannot reject the null hypothesis, indicating the existence of a unit root. The KPSS test gives results consistent with the ADF and PP tests, which reject the null hypothesis of no unit root. Thus all three test statistics indicate the existence of a unit root. Panel A of Table 10.7 also gives the unit root test result for the first difference of the price series. Note that neither the time trend nor an intercept is included in the test for the difference of the price series. The test results indicate the stationarity of the two variables.

Since the unit root test indicates that both of the closing price series have a unit root, we conduct the Johansen cointegration test and the result is shown in panel B of Table 10.7. Both the trace statistic and max-eigen statistic suggest that there is one cointegrating equation at the 5% significance level, which means that the logarithm prices of soybean futures in the DCE and the CBOT markets have similar stochastic trends and could form a long-run equilibrium relationship.

Table 10.8 shows the estimated result of the VEC model. As we can see from panel A of Table 10.8, the parameter of the cointegration equation is significant, and the sign of the estimated coefficient indicates that there is a positive long-term relationship between the closing price of the CBOT and the DCE. Panel B of Table 10.8 shows the estimated coefficient for the error correction term. The result suggests that both the DCE and the CBOT closing prices adjust with the long-term cointegration relationship, and the short-term adjustment is significant for both markets. Specifically, the significantly negative k_1 implies that a larger CN_CLOSE_{t-1} or a smaller US_CLOSE_{t-1} relative to the long-run equilibrium will result in a negative adjustment on CN_CLOSE_t , i.e. CN_CLOSE_t tends to revert to the long-run cointegration relationship. Similarly, the significantly positive k_2 implies that a larger US_CLOSE_{t-1} or a smaller CN_CLOSE_{t-1} relative to the long-run equilibrium will result in a negative adjustment on US_CLOSE_t , i.e. US_CLOSE_t tends to revert to the long-run cointegration relationship. We can thus indicate that the price from either the DCE or the CBOT tends to adjust to the price of the other exchange, i.e. the prices on both exchanges tend to influence each other.

Nonetheless, the error correction term in the equation with ΔCN_CLOSE_t being the dependent variable is greater in magnitude than that of the equation with ΔUS_CLOSE_t as the dependent variable. This

TABLE 10.7 Test Results of the Unit Root Test and the Cointegration Test. Panel A Gives the ADF, PP and KPSS Tests of the CBOT and the DCE Soybean Futures Price, as well as their First Difference. The Value in Parentheses is the Corresponding P-Value. Panel B Gives the Cointegration Test. Here ** Denotes Significant at the 5% Significance Level. The Sample Period is from March 15, 2002 to September 2, 2011 and the Series are the Log of the Original Price Series

Test statistic	CN_CLOSE	US_CLOSE	ΔCN_CLOSE	ΔUS_CLOSE
ADF	-1.53 (0.5187)	-1.25 (0.6556)	-32.93 (0)	-48.46 (0.0001)
PP	-1.55 (0.5063)	-1.32 (0.6243)	-51.95 (0.0001)	-48.53 (0.0001)
KPSS	4.448	4.117	0.101	0.051
Panel B: Cointegration test.				
Hypothesized No. of CE(s)	Trace statistic	Max-eigen statistic	5% Critical value	
			Trace statistic	Max-eigen statistic
None*	23.80**	19.55**	20.26	15.89
At most 1	4.25	4.25	9.16	9.16

Note: For the KPSS test, all the series are tested including the intercept; critical values at the 10%, 5%, and 1% significance level are 0.347, 0.463, and 0.7439, respectively.

TABLE 10.8 Estimated Results of the VEC Model. Panel A Gives the Cointegration Relationship between the Closing Price of the CBOT and the DCE Soybean Futures Price, While Panel B Gives the Estimated Coefficients in the Error Correction Equation. Here, *** Denotes the 1% Significance Level. The Sample Period is from March 15, 2002 to September 2, 2011 and the Series are the Log of the Original Price Series

CN_CLOSE_{t-1}	ΔUS_CLOSE_{t-1}	Constant
Panel A: The cointegration equation.		
1	-0.926*** [-16.53]	0.139
Panel B: The error correction equation.		
Independent variable	Dependent variable	
	ΔCN_CLOSE_t	ΔUS_CLOSE_t
Cointegration equation	-0.011*** [-3.68]	0.008* [1.95]
ΔCN_CLOSE_{t-1}	-0.130*** [-6.23]	0.007 [0.23]
ΔCN_CLOSE_{t-2}	-0.054*** [-2.70]	0.017 [0.56]
ΔUS_CLOSE_{t-1}	0.210*** [14.12]	-0.001 [-0.02]
ΔUS_CLOSE_{t-2}	0.030** [2.02]	0.008 [0.38]
Constant	-0.001 [-0.48]	0.001 [0.14]
Mon	0.003*** [4.13]	-0.001 [-0.03]

indicates that when the long-term equilibrium relationship is perturbed, it is the DCE price that makes the greater adjustment in order to re-establish the equilibrium. In other words, the CBOT price leads the DCE price in price discovery (Tse 1999).

10.4.3 Impulse Response and FEVD

We also carry out an impulse response analysis and forecast the error variance decomposition for the VEC model. The result for the impulse response analysis is shown in [Figure 10.6](#). As demonstrated in the figure, the CBOT closing price responds positively to an innovation in the DCE closing price, and the response is persistent. The DCE closing price also has a significant positive response to an innovation in the CBOT closing price, but the pattern of the response is different from that of the CBOT to the DCE shock. The response of the CBOT to the DCE shock immediately reaches a high level and then continues to increase gradually; however, the response of the DCE to the CBOT is almost 0 in the first period and then jumps to a high level in the second period, which persists for a long time. Nonetheless, the magnitude of the two responses to the other market is similar.

Using the method of Hasbrouck (1995), we analyse the information share attributable to each market with the forecasted error variance decomposition. An upper bound for the DCE's information share can be obtained by placing the DCE's price first in the Cholesky decomposition, while a lower bound may be obtained by placing its price last. The same reasoning applies to the CBOT. The results are presented in [Table 10.9](#) with the ordering of CN_CLOSE US_CLOSE on the left and the ordering of US_CLOSE CN_CLOSE on the right. The result indicates that about 0.1–5% of the forecasted error from the CBOT can be attributed to the DCE, while about 5–25% of the forecasted error from the DCE can be attributed to the CBOT. It is clear from the result that although information flows bidirectionally in the two markets, the CBOT still leads the DCE.




FIGURE 10.6 The impulse response analysis in the VEC model for the closing price of the DCE and the CBOT soybean futures. The sample period is March 15, 2002 to September 2, 2011, and the series are the log of the original price series.

TABLE 10.9 Forecasted Error Variance Decomposition of the DCE and CBOT Closing Price of the Soybean Futures in the VEC Model. Panel A is the Variance Decomposition for the DCE Closing Price, While Panel B is for the Cbot Closing Price. In Each Panel, the First Column is the Forecast Period, The Second and the Third Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_CLOSE* and *US_CLOSE* With the Cholesky Ordering of *CN_CLOSE* *US_CLOSE*, While The Fourth and Fifth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_CLOSE* and *US_CLOSE* with the Cholesky Ordering of *US_CLOSE* *CN_CLOSE*. The Sample Period is March 15, 2002 to September 2, 2011 and the Series are the Log of the Original Price Series

Period	<i>CN_CLOSE</i>	<i>US_CLOSE</i>	<i>CN_CLOSE</i>	<i>US_CLOSE</i>
Panel A: Variance decomposition of <i>CN_CLOSE</i>				
1	100.00	0.00	96.57	3.43
2	95.02	4.98	86.58	13.42
3	93.35	6.65	83.09	16.91
4	92.09	7.91	80.78	19.22
5	91.26	8.74	79.31	20.69
6	90.59	9.41	78.18	21.82
7	90.04	9.96	77.27	22.73
8	89.55	10.45	76.50	23.50
9	89.11	10.89	75.81	24.19
10	88.69	11.31	75.18	24.82

(Continued)

TABLE 10.9 (Continued) Forecasted Error Variance Decomposition of the DCE and CBOT Closing Price of the Soybean Futures in the VEC Model. Panel A is the Variance Decomposition for the DCE Closing Price, While Panel B is for the Cbot Closing Price. In Each Panel, the First Column is the Forecast Period, The Second and the Third Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_CLOSE* and *US_CLOSE* With the Cholesky Ordering of *CN_CLOSE US_CLOSE*, While The Fourth and Fifth Columns Correspond to the Percentage of the Forecasted Error that can be Attributed to *CN_CLOSE* and *US_CLOSE* with the Cholesky Ordering of *US_CLOSE CN_CLOSE*. The Sample Period is March 15, 2002 to September 2, 2011 and the Series are the Log of the Original Price Series

Period	<i>CN_CLOSE</i>	<i>US_CLOSE</i>	<i>CN_CLOSE</i>	<i>US_CLOSE</i>
Panel B: Variance decomposition of <i>US_CLOSE</i>				
1	3.43	96.57	0.00	100.00
2	3.62	96.38	0.01	99.99
3	3.89	96.11	0.03	99.97
4	4.06	95.94	0.04	99.96
5	4.21	95.79	0.06	99.94
6	4.34	95.66	0.07	99.93
7	4.47	95.53	0.09	99.91
8	4.59	95.41	0.11	99.89
9	4.71	95.29	0.13	99.87
10	4.82	95.18	0.15	99.85
Cholesky ordering: <i>CN_CLOSE US_CLOSE</i>			Cholesky ordering: <i>US_CLOSE CN_CLOSE</i>	

10.5 Conclusion

In this paper, we examine the role that the DCE plays in the price discovery of soybean futures. We employ the SVAR and VEC models to investigate the information transfer between soybean futures prices on the DCE and the CBOT. In addition to the well-documented fact that the CBOT significantly affects the DCE, our analysis also emphasizes that the DCE simultaneously has a significant impact on the CBOT. More importantly, the magnitude of the impacts is similar in both directions from the results of the forecasted error variance decomposition (information share) in the SVAR and VEC models. This result logically leads us to the conclusion that the DCE is already playing an important role in the global price discovery of soybean futures, and the information revealed in its trading can affect global market pricing. Because China is now the largest soybean importer in the world and the trading volume of soybean futures on the DCE is second only to the CBOT, it is reasonable that the Chinese market is playing a prominent role in the global price discovery of soybean futures.

Acknowledgements

We thank Shang Liu for his excellent research assistance with preliminary data analysis. Han acknowledges the support from the National Natural Sciences Foundation of China (NSFC), project No. 70831001. Tang acknowledges financial support from the National Natural Science Foundation of China (grant No. 71171194).

References

- Booth, G.G. and Ciner, C., International transmission of information in corn futures markets. *J. Multinatl Financ. Mgmt*, 1997, 7, 175–187.
- Eun, C.S. and Sabherwal, S., Cross-border listings and price discovery: evidence from U.S.-listed Canadian stocks. *J. Finance*, 2003, 58, 549–575.

- Fung, H.G., Leung, W.K. and Xu, X.E., Information flows between the U.S. and China commodity futures trading. *Rev. Quantit. Finance Account*, 2003, **21**, 267–285.
- Garbade, K.D. and Silber, W.L., Dominant and satellite markets: a study of dually-traded securities. *Rev. Econ. Statist.*, 1979, **61**, 455–460.
- Grammig, J., Melvin, M. and Schlag, C., Internationally cross-listed stock prices during overlapping trading hours: price discovery and exchange rate effects. *J. Empir. Finance*, 2005, **12**, 139–164.
- Harris, F.H., McNish, T.H., Shoesmith, G.L. and Wood, R.A., Cointegration, error correction, and price discovery on informationally linked security markets. *J. Financ. Quantit. Anal.*, 1995, **30**, 563–579.
- Hasbrouck, J., One security, many markets: determining the contributions to price discovery. *J. Finance*, 1995, **50**, 1175–1199.
- Kleidon, A.W. and Werner, I.M., U.K. and U.S. trading of British cross-listed stocks: an intraday analysis of market integration. *Rev. Financ. Stud.*, 1996, **9**, 619–664.
- Tse, Y., Price discovery and volatility spillovers in the DJIA index and futures markets. *J. Fut. Mkts*, 1999, **19**(8), 911–930.
- Tse, Y. and Booth, G.G., Information shares in international oil futures markets. *Int. Rev. Econ. Finance*, 1997, **6**, 49–56.
- Tse, Y., Lee, T.-H. and Booth, G.G., The international transmission of information in Eurodollar futures markets: a continuously trading market hypothesis. *J. Int. Money Finance*, 1996, **15**, 447–465.
- Xu, X.E. and Fung, H.G., *J. Int. Financ. Mkts, Inst. Money*, 2005, **15**, 107–124.

The Structure of Gold and Silver Spread Returns

Jonathan A. Batten	11.1 Introduction	213
Cetin Ciner	11.2 Data	215
Brian M. Lucey	11.3 Rescaled Adjusted Range Analysis	217
Peter G. Szilagyi	11.4 Trading Implications.....	219
	11.5 Conclusions.....	223
	References.....	223

The price dynamics of gold and silver have long been a matter of popular concern and fascination. The objective of this study is to investigate the dynamics of the bivariate relationship between gold and silver prices. First, we investigate the spread, measured as the price difference between gold and silver trading as a futures contract. Then the presence of a fractal structure is measured using statistical techniques based on rescaled range analysis after accommodating short-term autocorrelated innovations in the return process. To highlight the economic consequences of fractality, we apply trading rules based upon the Hurst coefficient to the time series data. Importantly, we find that these rules outperform simple buy-hold and moving-average strategies over varying holding periods.

Keywords: Anomalies in prices; Chaos theory; Commodity prices; Econophysics; Risk management

JEL Classification: C1, C19, F3, F30, G1, G10

11.1 Introduction

The price dynamics of precious metals, generally, and gold and silver, in particular, has long been a matter of popular concern and fascination. Recently, a number of authors have successfully modelled the stochastic nature of precious metal returns (e.g., Urich (2000), Casassus and Collin-Dufresne (2005), Cochran *et al.* (2012) and Baur (2013)), while new tools from econometrics have demonstrated important insights into the long-term price relationships between bivariate combinations of precious metal prices (e.g., Escibano and Granger (1998) and Ciner (2001)).

The objective of this study is to report the long-run price dynamics of the bivariate relationship between gold and silver. First, we investigate the spread, measured as the price difference between gold and silver trading as a futures contract. This is novel since previous analyses tend to focus on the price of individual series. Then, we extend the work of Ciner (2001) and Figuerola-Ferretti and Gonzalo (2010), who apply the cointegration techniques of Johansen (1991) and Escibano and Granger (1998) to identify the presence of a long-term equilibrium in the gold and silver markets, to the investigation of the presence of long-term dependence—or long memory—effects in the time-series of the gold-silver spreads.

The investigation and subsequent identification of complex dynamics, including long-memory processes, in various financial time series has provoked debate in the empirical finance literature (e.g., Mandelbrot (2001) and Bianchi and Pianese (2007)), especially in terms of how the presence of these processes affect asset market efficiency (e.g., Eom *et al.* (2008) and McCauley *et al.* (2008)) and asset pricing (e.g., Ellis and Hudson (2007) and Takami *et al.* (2008)). Furthermore, a key contribution of this paper is that we are able to assess the economic implications of the identified long-memory process by applying a series of trading rules to the gold-silver spread time-series. Trading rules may be used to test the efficiency of markets (e.g., Brock *et al.* (1992)), since no trading rule should consistently be able to outperform a simple buy-hold strategy in the long term.

Long-memory processes are typically associated with the hyperbolic decay of the autocorrelation function and are easily measured using the classical rescaled adjusted range (RAR) technique of Hurst (1951), although other approaches, such as detrended fluctuation analysis (e.g., Wan *et al.*'s (2011a,b) analysis of the gold and oil markets respectively) may be used.* One key advantage of the RAR approach is that, given a daily time series of length N , the RAR may be estimated over a rolling sample (n) to produce a series of (daily) statistics (of length $N - n + 1$). In our case, we set $n = 22$ and 66 , representing 1 month and 3 months, respectively, and utilise a time series of Hurst statistics of 1682 daily observations. The RAR approach also has the benefit of providing an insight into the direction of the equilibrium reverting process since it allows for differentiation between processes that revert to their long-term mean after an information shock, such as from news announcements (Christie-David *et al.* 2000), and those that progressively move away from the long-term mean after each new shock.

One feature common to most financial time series is the presence of short-memory effects (e.g., Lo (1991)), typically observed as short-term autocorrelation, which may be associated with lingering liquidity effects in a financial market. To overcome this feature of the gold-silver spread returns, we filter the series using AutoRegressive Moving Average (ARMA) techniques prior to estimating the RAR. This approach has been used previously by Szilagyi and Batten (2007) and Batten and Hamada (2009) to prefilter residuals from foreign exchange and electricity returns before testing for fractality using the rescaled range approach. The ARMA filtering approach therefore accommodates any short-term autocorrelated innovations that may be present in the return process.

We make two main contributions in this paper. First, the gold-silver spread returns reveal time-varying non-linearities, which can be identified as fractality using commonly applied Hurst tests. The use of the local Hurst coefficient, estimated over a rolling sample period (or window) of 22-day and 66-day periods, highlights the time-varying nature of this phenomenon. Time dependence in the Hurst coefficient has been identified by a number of empirical studies investigating the dynamics of financial prices, including Carbone *et al.* (2004), Batten *et al.* (2008) and Grech and Pamela (2008), amongst others, and so we add to this literature. Second, we test the performance of simple trading rules based upon the Hurst coefficient and, importantly, find that these rules outperform both a simple moving-average strategy (which is typically applied to trending series by traders) and a buy-hold strategy. This finding adds to a developing literature that utilises the Hurst coefficient as an investment and trading tool (e.g., Clark (2005)).

These findings are also consistent with other researchers that identify time-varying long-term dependence in other financial markets (Lo 1991, Batten *et al.* 2005, Cajueiro and Tabak 2007, 2008, Du and Ning 2008, Grech and Pamula 2008), other forms of non-linearities or complex structure in financial returns (Jiang and Zhou 2008, Takami *et al.* 2008), or demonstrated the economic gains from trading strategies designed to exploit stock market inefficiencies—specifically trend following trading systems (Brock *et al.* 1992, Bessembinder and Chan 1998, Eom *et al.* 2008), or the possible effects of thin and non-synchronous trading, typical of emerging markets (Bley 2011).

* Owczarczuk (2012) provides an excellent discussion of this technique. The links between oil and other commodity markets is discussed in Mohanty *et al.* (2010) and Ewing and Malik (2013).

Next we briefly provide information on the gold and silver futures data used for the study. Then the results from price analysis using rescaled range analysis are reported and the trading strategy applied. The final section allows for some concluding remarks.

11.2 Data

Lucey and Tulley (2006) provide a detailed account of studies investigating trading in the international gold and silver markets. Gold is an important reserve asset and in recent times has become an important part of the monetary regime in emerging economies (see Taguchi 2011). We investigate the price of two contracts trading on the New York Mercantile Exchange (NYMEX): the deliverable 100 troy ounce nominal COMEX gold and the deliverable 5000 troy ounce nominal COMEX silver contracts. In an economic sense these futures contract are fully arbitrageable against gold and silver trading in a variety of other worldwide cash and futures markets. Open-outcry trading commences at 08:20 h/08:25h and ends at 13:30h/13:25 h (gold/silver). Trading is also available simultaneously (termed side-by-side trading) on the GLOBEX electronic trading system available on the Chicago Mercantile Exchange (CME). Our data comprises 1746 daily observations of the near month COMEX gold and silver contract at the start of trading from January 1999 to December 2005. This number is reduced to 1682 once the first 64 observations are used to begin the rolling Hurst estimations.

We first estimate the interday returns (ΔP_t) for the price spread (P_t) between gold (G_t) and silver (S_t), where $P_t = G_t - S_t$. Allow $\Delta P_t = \log(P_t) - \log(P_{t-1})$ where the interval $t - 1 \rightarrow t$, is 1 day. Individual asset returns (for gold and silver) are also measured as $\Delta G_t = \log(G_t) - \log(G_{t-1})$ and $\Delta S_t = \log(S_t) - \log(S_{t-1})$ where the interval $t - 1 \rightarrow t$, is 1 day.

Note that to the extent the returns of the underlying assets (ΔG_t and ΔS_t) are themselves random processes, then $\Delta P_t = \varepsilon_t$, with ε_t being a random variable, which is expected to have a mean of zero. ΔP_t should also be both mean stationary and uncorrelated over various time increments, which is a requirement for efficient markets in the sense of Fama (1998), although, as is well known, these features are rarely present in financial markets (McCauley *et al.* 2008).

The descriptive statistics of our data are reported in Table 11.1. Over the sample period the spread varied from a maximum of US\$474.31 to a minimum of US\$248.63. The spread mean was US\$329.32. For the return series (Figures 11.1(d) and 11.1(e)) the mean was slightly positive in all cases (ΔP_t , ΔG_t , ΔS_t equalling 0.00013, 0.00014 and 0.00013, respectively) with variation on a similar scale as shown in

TABLE 11.1 Descriptive Statistics of the Gold–Silver Spread (Daily Spreads From January 1999 to December 2005)

	Gold–silver spread $P_t = G_t - S_t$	Gold–silver spread return $\Delta P_t = \log(P_t) - \log(P_{t-1})$
Mean	329.32	0.00013
Median	309.44	0.00015
Maximum	474.31	0.04464
Minimum	248.63	−0.02873
Standard deviation	62.55	0.00440
Skewness	0.505	0.902
Kurtosis	1 801	16.411
Jarque–Bera	172.34	12 839.13
Probability	0.000	0.00000
Observations	1682	1682

The original sample contains 1746 observations. The later 22-day (66-day) estimation of the Hurst coefficient is based on the returns from $R_0 \rightarrow R_{t-22}$ (and $R_0 \rightarrow R_{t-66}$) observations. To allow convergence to a stable Hurst coefficient we simply report the subsequent 1682 observations.




FIGURE 11.1 (a) The coefficient of variation of daily silver returns estimated on a 22-day rolling window. (b) The coefficient of variation of daily gold returns estimated on a 22-day rolling window. (c) The correlation between daily gold and silver returns estimated on a 22-day rolling window.

Figures 11.1(d) and 11.1(e). Silver was more volatile over the entire sample period than the spread, or gold, measured by the standard deviation (ΔP_p , ΔG_p , ΔS_t equalling 0.0044, 0.0044 and 0.0062, respectively) and the coefficient of variation (ΔP_p , ΔG_p , ΔS_t equalling 3059, 3060 and 4733, respectively). However, this was not the case when the coefficient of variation (CV) was estimated over a 22-day (equal to one calendar month) rolling window (i.e. $CV_n = \mu_n / \sigma_n$, where $n = 22$). The CV_{22} for gold and silver are illustrated in Figures 11.1(a) and 11.1(b). It is clear from these figures that silver appears more stable, with less significant peaks and troughs in the CV_{22} of returns than gold. Such differences would likely impact upon the dynamics of the long-term equilibrium between the two metals.




FIGURE 11.1 (Continued) (d) The average 22-day daily gold and silver returns estimated on a rolling window. (e) The average 22-day gold–silver spread returns estimated on a rolling window.

To demonstrate the time-varying nature of the relationship between gold and silver we also estimate and then plot in [Figure 11.1\(c\)](#) the rolling 22-day correlation between gold and silver returns (GSp_{22}). Over the entire sample period the correlation is high and positive ($GSp_{1746} = 0.686$, $p = 0.000$). However, estimating GSp_{22} the range varies from a maximum of 0.9675 to a minimum of -0.1524 . Nonetheless, since a positive correlation is maintained over the sample period, trading strategies based on mean reversion of the spread to its average may in fact provide profitable opportunities for market participants. This finding is consistent with widely held views in commodity markets that despite the fundamental differences between the two markets, gold and silver prices tend to move together (Lucey and Tulley 2006), thereby offering the possibility for various trading and portfolio strategies that exploit mean reversion in the spread returns. This possibility is considered further when the rescaled range statistic is estimated in the next section and when trading strategies based upon the local Hurst exponent are implemented in [Section 11.4](#).

11.3 Rescaled Adjusted Range Analysis

The presence of long-term dependence in the spread returns (ΔP_t) between gold and silver may be measured using statistical techniques based on the Hurst (1951) rescaled range analysis, after accommodating for any short-term autocorrelated innovations in the return process (Batten *et al.* 2008). The

filtering process is readily accomplished via ARMA models and forms the basis for RAR (Szilagyi and Batten 2007, Batten and Hamada 2009). Of specific interest is the residual ψ_t after applying various filters (AR(0) $\rightarrow$ ARMA(2,1)) to ΔP_t . Consider an ARMA(2,1) model of the form

$$\Delta P_t = \alpha_0 + \beta_1 \Delta P_{t-1} + \beta_2 \Delta P_{t-2} + X_1 \lambda \psi_{t-1} + \psi_t, \quad (11.1)$$

which systematic analysis is found to provide the best fit to the data with $\beta_1 = 0.4802$ ($p = 0.06$), $\beta_2 = 0.1109$ ($p = 0.000$) and $X_1 = 0.5780$ ($p = 0.024$).

For each ψ_t over the subsample n , the classical rescaled adjusted range $(R/\sigma)_n$ of Hurst (1951) and Mandelbrot and Wallis (1969) is calculated as

$$\left(\frac{R}{\sigma}\right)_n = \frac{1}{\sigma_n} \left[\max_{1 \leq k \leq n} \sum_{j=1}^k (\psi_j - \mu_n^\psi) - \min_{1 \leq k \leq n} \sum_{j=1}^k (\psi_j - \mu_n^\psi) \right], \quad (11.2)$$

where μ_n is the mean and σ_n is the standard deviation of ψ_t over an overlapping sample of length n

$$\sigma_n = \left[1/n \sum_{j=1}^n (\psi_j - \mu_n^\psi)^2 \right]^{0.5}. \quad (11.3)$$

In order to capture the time-varying nature of the dependence in ψ_t this study employs a local measure of the Hurst exponent (h) as used by Szilagyi and Batten (2007) and Batten and Hamada (2009). The local version (h) of the Hurst exponent is then estimated for $(N - n + 1)$ times overlapping subseries of these various length n :

$$h_n = \frac{\log(R/\sigma)_n}{\log n}, \quad (11.4)$$

where n is set to either 22 days or 66 days, which is equivalent to a standard one- and three-month period. This procedure in effect creates a time-series of exponent values,* the change in whose value can be measured over time. The averages of the local Hurst (h) for the entire sample period are summarised in Table 11.2. The top row in this table records the filtering technique applied to ψ_t . These four techniques range from AR(0)—no filtering—to ARMA(2, 1) as per Equation (11.1).

TABLE 11.2 Local Hurst Exponents Estimated Using 22- and 66-Day Rolling Windows

	22-Day rolling window (1 month)				66-Day rolling window (3 months)			
	AR(0)	AR(1)	AR(2)	ARMA(2,1)	AR(0)	AR(1)	AR(2)	ARMA(2,1)
Overall sample								
Mean	0.7071	0.7370	0.7189	0.7170	0.6487	0.6667	0.6577	0.6496
Standard deviation	0.1314	0.1328	0.1312	0.1283	0.0737	0.0741	0.0740	0.0725
95% Confidence interval	0.7007–0.7133	0.7239–0.7366	0.7123–0.7521	0.7046–0.7168	0.6451–0.6523	0.6631–0.6703	0.6541–0.6613	0.6460–0.6532

* The filtering using the ARMA(2,1) affects the residuals of Equation (11.1). It is the residual series ψ_t that is used to estimate the $(R/\sigma)_n$ of Hurst (1951) using $(R/\sigma)_n = 1/\sigma_n$. The transformation shown in Equation (11.4) is then applied to estimate the local Hurst (h) over a 22-day and 66-day rolling window.

Recall from Hurst (1951) that, under the null hypothesis of no long-term dependence, the value of $h_n = 0.5$ (a white noise process). For time-series exhibiting positive long-term dependence, the observed value of the exponent is $h_n > 0.5$. Time-series containing negative dependence are mean-reverting and alternatively characterised by $h_n < 0.5$ (termed pink noise by Mulligan (2004)). Note that while the original sample contains 1746 observations, estimation of the 22-day (66-day) local Hurst coefficient is based on the returns from $R_0 \rightarrow R_{t-22}$ (and $R_0 \rightarrow R_{t-66}$) observations. To allow convergence to a stable Hurst coefficient we ignore the first 64 estimations from the original sample.

From Table 11.1, the mean for h_{22} varies from 0.7070 for AR(0) to 0.7170 for ARMA(2,1). The scores for h_{66} are lower, suggesting the series is becoming more random as the sample length n in Equation (11.4) increases, and vary from 0.6487 for AR(0) to 0.6496 for ARMA(2,1). Note also that these different filters have little effect on the size of the h -statistic as evidenced by the overlapping of the confidence intervals. This is due to the fact that the long memory, by definition, relates to the lingering effects of hyperbolic decay in the autocorrelation of the return series. One economic explanation for this property is due to liquidity effects in the market. Even allowing for a 95% confidence interval, these local Hurst coefficients are consistent with positive long-term dependence. For positively dependent processes, another price movement further away from the mean (or long-term equilibrium) will follow the earlier movement away from equilibrium. Thus, they are trend-reinforcing processes (Mulligan 2004), such that the spread should tend to become larger, or smaller, depending on whether the previous change in price was positive, or negative.

Nonetheless, a plot of h_{22} and h_{66} over the entire sample period (shown in Figures 11.2(a) and 11.2(b)) shows considerable variation in the statistics and rare episodes when the statistic was below 0.5000. In the case of h_{22} the minimum value was 0.3296 and the maximum was 1.1667, while for h_{66} the minimum value was 0.4071 and the maximum was 0.9399. Values below 0.5000 are consistent with negative dependence where a price movement towards the equilibrium should follow a movement away from equilibrium. Note that this high level of the Hurst exponent (greater than 1) appears to be a statistical artefact due to extreme volatility in the daily gold–silver spread price which moved from USD\$250 to US\$320 in the period from 9/7/1999 to 11/3/1999 (nearly a 30% move). Removing the spreads from the period 10/1/1999 to 10/18/1999 eliminates this anomaly. Alternately, if the estimation sample length is increased to 66, then the anomaly dissipates. Removing these extreme observations does affect the estimated average Hurst value. For example, in Table 11.2 the ARMA(2,1) 22-day Hurst is now 0.7048 versus the reported number of 0.7170 and, for the ARMA(2,1), the 66-day Hurst is 0.6475 versus the reported number of 0.6496, which lie just within the confidence interval.

11.4 Trading Implications

To investigate the trading implications we first divide the unfiltered return series (ΔP_t) into varying holding periods (HP) from 1 day to 22 days from the start day return at t_0 (i.e. R_0). These holding period returns are reported in Table 11.3 and reflect the reward to the investor for a buy-hold investment strategy for these specified number of days. Note that kurtosis and skewness decline with holding period length, while the mean return increases as days are added to the holding period. However, the return over the holding period on an average daily basis is not linear. This is clearer in Figure 11.3, which plots the average daily return over the varying holding period lengths (reported in the third column of Table 11.3). Figure 11.3 clearly shows that there is an advantage to holding a portfolio for the next 2 days instead of just one day, but this advantage declines quickly for the next 6 days, after which it increases again steadily for the next period. The optimal holding period (that is the holding period with the largest average daily return) is a period of 15 days from today.

A number of trading strategies were then applied to the return series. The first of these is one commonly applied to trend-following systems: a moving average. Brock *et al.* (1992), in a leading paper, demonstrated the success of moving average trading systems over others. While moving averages of varying





FIGURE 11.2 (a) The local Hurst estimated using an ARMA(2,1) filter on the spread return between daily gold and silver (estimated on a 22-day rolling window). (b) The local Hurst estimated using an ARMA(2,1) filter on the spread return between daily gold and silver (estimated on a 66-day rolling window).

length could be considered, for the sake of brevity we investigate a moving average of length 22, which matches the other estimations for the local Hurst coefficient.

The first rule that is applied is to buy when the 22-day moving average price is greater than the price today (i.e. $P_t > \text{MA}_{22}$ in Table 11.4). If this rule is triggered, the subsequent returns to the next day ($\text{HP} = R_0 \rightarrow R_1$), the next 5 days ($\text{HP} = R_0 \rightarrow R_5$), the next 10 days ($\text{HP} = R_0 \rightarrow R_{10}$) and the next 20 days ($\text{HP} = R_0 \rightarrow R_{20}$) are reported in the first four rows of Table 11.4. The frequency of this rule is reported as $N = 0$ (reject) and $N = 1$ (accept). In all cases there were more accepted applications of the rule than rejected. However, while mean returns from the accepted application of the rule ($N = 1$) steadily increased from 0.00133 for the next day to 0.00252 for the next 20 days, an F -test of difference in the means shows that only the next one-day and 5-day holding periods were statistically different from the alternate (rejected) portfolio returns ($N = 0$). Importantly, the returns for the return on the average buy-hold portfolio held for 20 days (reported in Table 11.3 as 0.003079) is higher than the return for an accepted application of the rule for 20 days (Table 11.4, $N = 1$, mean 0.00252). Thus, while the moving

TABLE 11.3 Returns Over Different Holding Periods (1 Day to 22 Days)

Return (R) window	Mean	Average daily	Standard deviation	Kurtosis	Skewness
R_0	0.0001443	0.000144349	0.004417358	0.8478223	12.916320
$R_0 \rightarrow R_1$	0.0002910	0.000145520	0.005920985	1.0070222	8.732371
$R_0 \rightarrow R_2$	0.0004361	0.000145371	0.007283346	1.0934505	9.641540
$R_0 \rightarrow R_3$	0.0005775	0.000144373	0.008408646	0.9301216	6.773518
$R_0 \rightarrow R_4$	0.0007179	0.000143574	0.009547577	1.0231929	6.256877
$R_0 \rightarrow R_5$	0.0008565	0.000142749	0.010560701	1.0968331	7.211523
$R_0 \rightarrow R_6$	0.0009934	0.000141918	0.011531883	1.1706315	8.118051
$R_0 \rightarrow R_7$	0.0011374	0.000142175	0.012337697	1.1952264	7.778539
$R_0 \rightarrow R_8$	0.0012822	0.000142462	0.013090957	1.1863138	7.273252
$R_0 \rightarrow R_9$	0.0014258	0.000142576	0.013770619	1.2489172	7.404415
$R_0 \rightarrow R_{10}$	0.0015711	0.000142826	0.014435979	1.2597065	6.978512
$R_0 \rightarrow R_{11}$	0.0017207	0.000143395	0.015047962	1.2411078	6.672399
$R_0 \rightarrow R_{12}$	0.0018772	0.000144402	0.015588680	1.2269747	6.462010
$R_0 \rightarrow R_{13}$	0.0020384	0.000145601	0.016069123	1.2125824	6.145907
$R_0 \rightarrow R_{14}$	0.0021926	0.000146175	0.016539234	1.2111407	5.949431
$R_0 \rightarrow R_{15}$	0.0023446	0.000146540	0.016999905	1.1954588	5.660918
$R_0 \rightarrow R_{16}$	0.0024945	0.000146736	0.017448983	1.1795708	5.321579
$R_0 \rightarrow R_{17}$	0.0026419	0.000146773	0.017885264	1.1508532	4.945140
$R_0 \rightarrow R_{18}$	0.0027895	0.000146814	0.018343986	1.0978180	4.601016
$R_0 \rightarrow R_{19}$	0.0029360	0.000146802	0.018774185	1.0558345	4.244551
$R_0 \rightarrow R_{20}$	0.0030790	0.000146621	0.019170964	1.0150762	3.913951
$R_0 \rightarrow R_{21}$	0.0032216	0.000146438	0.019545214	0.9705729	3.697957
$R_0 \rightarrow R_{22}$	0.0033633	0.000146233	0.019907824	0.9214373	3.398959

**FIGURE 11.3** Plot of average daily gold-silver spread returns for holding periods from one day to one month.

average rule offers potential short-term gains, these appear to diminish as the holding period is lengthened. This short-term gain is consistent with trading models that exploit the autoregressive features of the series. These were previously identified in the ARMA model (Equation (11.1)), which identifies positive one- and two-day autocorrelation in the return series.

TABLE 11.4 Analysis of Variance (Anova) of Differences in Mean Returns Based on Various Trading Rules Over Different Holding Periods.

Trading rule	Mean				Standard deviation		<i>F</i> -Test	<i>p</i> -Value	Adjusted <i>R</i> ² (%)
	<i>N</i> = 0	<i>N</i> = 1	<i>N</i> = 0	<i>N</i> = 1	<i>N</i> = 0	<i>N</i> = 1			
$P_t > \text{MA22}$ $\text{HP} = R_0 \rightarrow R_1$	812	889	-0.00085	0.00133	0.004834	0.00665	58.59	0.000	3.28
$P_t > \text{MA22}$ $\text{HP} = R_0 \rightarrow R_5$	810	887	-0.00001	0.00164	0.00846	0.01221	10.18	0.001	0.54
$P_t > \text{MA22}$ $\text{HP} = R_0 \rightarrow R_{10}$	808	884	0.00108	0.00206	0.01265	0.01603	1.96	0.162	0.06
$P_t > \text{MA22}$ $\text{HP} = R_0 \rightarrow R_{20}$	808	874	0.00387	0.00252	0.01979	0.01876	2.08	0.150	0.06
$P_p \text{H22} < 0.4$ $\text{HP} = R_0 \rightarrow R_1$	1,674	27	0.00031	-0.00073	0.00591	0.00829	0.81	0.368	0.00
$P_p \text{H22} < 0.4$ $\text{HP} = R_0 \rightarrow R_5$	1,670	27	0.00098	-0.00681	0.01061	0.00775	14.40	0.000	0.78
$P_p \text{H22} < 0.4$ $\text{HP} = R_0 \rightarrow R_{10}$	1,665	27	0.00181	-0.01184	0.01449	0.00836	23.81	0.000	1.33
$P_p \text{H22} < 0.4$ $\text{HP} = R_0 \rightarrow R_{20}$	1,655	27	0.00345	-0.01425	0.01925	0.01072	22.71	0.000	1.28
$P_p \text{H22} > 0.6$ $\text{HP} = R_0 \rightarrow R_1$	359	1342	-0.00021	0.00042	0.00621	0.00588	3.14	0.076	0.13
$P_p \text{H22} > 0.6$ $\text{HP} = R_0 \rightarrow R_5$	359	1338	-0.00071	0.00127	0.01166	0.01028	9.96	0.002	0.53
$P_p \text{H22} > 0.6$ $\text{HP} = R_0 \rightarrow R_{10}$	359	1333	-0.00128	0.00237	0.01442	0.01445	18.05	0.000	1.00
$P_p \text{H22} > 0.6$ $\text{HP} = R_0 \rightarrow R_{20}$	359	1323	-0.00214	0.00461	0.01727	0.01953	35.36	0.000	2.00

The original sample contains 1746 observations. The 22-day Hurst coefficient is based on the previous 22 observations, which allows 1746–22 Hurst (*H*) estimations. To allow convergence to a stable Hurst coefficient we ignore the next 22 observations. The holding period (HP) return (*R*) varies from the return to the next day ($\text{HP} = R_0 \rightarrow R_1$) to the next 20 days ($\text{HP} = R_0 \rightarrow R_{20}$). Therefore, the sum of *N* = 0 and *N* = 1 varies from 1701 (for a 1-day holding period) to 1652 (for a 20-day holding period). The trading rules are buy when the price at *t* = 0 > previous 22-day Moving Average ($P_t > \text{MA22}$), sell the price at *t* = 0 when the previous 22-day Hurst coefficient is < 0.4 ($P_p \text{H22} < 0.4$), and buy the price at *t* = 0 when the previous 22-day Hurst coefficient is > 0.6 ($P_p \text{H22} > 0.6$).

The next set of rules that are applied utilise the time-series of local Hurst coefficients, whose averages are reported in Table 11.2. For the sake of brevity, only the unfiltered local Hurst (AR0) estimated over 22 days is reported, although similar findings apply for the Hurst (AR0) 66-day estimations. There are two conditions considered to provide the perspective when the return time series is mean reverting ($\text{H22} < 0.4$; 27 instances when the rule is accepted), or trend-reinforcing ($\text{H22} > 0.6$; 1342 instances when the rule is accepted). The same holding period lengths for the moving average rule are applied (next day, next 5 days, next 10 days and next 20 days). Apart from one exception (for $\text{H22} < 0.4$, for the next day holding period), which is the likely consequence of the positive autoregressive structure mentioned earlier, the *F*-test of the difference in the other mean returns shows significant differences in returns over the alternate portfolio. Also, these returns are significantly different to the returns from the average buy-hold portfolios of varying holding periods reported in Table 11.3. For example, from Table 11.4, for the Hurst rule of $\text{H22} > 0.6$, for the 20-day holding period, the mean return is 0.00461 compared with the

20-day buy-hold of 0.003079 (Table 11.2). Note that to effect a profitable trade the investor is required to sell when $H22 < 0.4$ (a bet on mean reversion), and buy when $H22 > 0.6$ (a bet on trend-reinforcement).

11.5 Conclusions

This study investigates the relationship between gold and silver trading as a futures contract on COMEX from January 1999 to December 2005. During this period the correlation between gold and silver returns was positive and high, even though the relationship itself was unstable. We apply techniques from fractal geometry after accommodating underlying autoregressive behaviour to investigate the long-term dynamics of the spread between these two contracts. Using a local Hurst exponent we find episodes of both positive and negative dependence, although the positive dependent relationship appears to be dominant. This last finding is suggestive of a time-varying fractal structure in the spread returns. Positive dependence (consistent with a Hurst coefficient > 0.5) in the gold–silver spread returns suggests the series will not immediately revert to its average or long-term mean, thereby offering traders limited profit opportunities. To test this proposition we test a number of simple trading rules based upon the Hurst coefficient.

We find that trading rules requiring the investor to sell when the local Hurst coefficient, estimated over a 22-day window, is less than 0.4 (a bet on mean reversion), and buying when the local Hurst coefficient, estimated over a 22-day window, is greater than 0.6 (a bet on trend-reinforcement) outperforms a simple buy-hold and moving-average strategy. This result adds to a growing body of work that finds significant arbitrage possibilities remaining in financial markets despite the advent of improved pricing and information technology (Avellaneda and Lee 2010).

References

- Avellaneda, M. and Lee, J.-H., Statistical arbitrage in the US equities market. *Quant. Finance*, 2010, **10**(7), 761–782.
- Batten, J.A., Ellis, C. and Fetherston, T., Statistical illusions and long-term return anomalies on the Nikkei. *Chaos, Solitons & Fractals*, 2005, **23**, 1125–1136.
- Batten, J.A., Ellis, C. and Fetherston, T., Sample period selection and long-term dependence: New evidence from the Dow Jones. *Chaos, Solitons & Fractals*, 2008, **36**(5), 1126–1140.
- Batten, J.A. and Hamada, M., The compass rose pattern in electricity prices. *Chaos: Interdiscip. J. Nonlin. Sci.*, 2009, **19**, 043106.
- Baur, D.G., The autumn effect of gold. *Res. Int. Bus. Finance*, 2013, **27**(1), 1–11.
- Bessembinder, H. and Chan, K., Market efficiency and the returns to technical analysis. *Financ. Mgmt.*, 1998, **27**(2), 5–17.
- Bianchi, S. and Pianese, A., Modelling stock price movements: Multifractality or multifractionality? *Quant. Finance*, 2007, **7**(3), 301–319.
- Bley, J., Are GCC stock markets predictable? *Emerg. Mkts Rev.*, 2011, **12**(3), 217–237.
- Bradley, T.E. and Malik, F., Volatility transmission between gold and oil futures under structural breaks. *Int. Rev. Econ. Finance*, 2013, **25**, 113–121.
- Brock, W., Lakonishok, J. and LeBaron, B., Simple technical trading rules and the stochastic properties of stock returns. *J. Finance*, 1992, **47**(5), 1731–1764.
- Cajueiro, D.O. and Tabak, B.M., Long-range dependence and market structure. *Chaos, Solitons & Fractals*, 2007, **31**(4), 995–1000.
- Cajueiro, D.O. and Tabak, B.M., Testing for long-range dependence in world stock markets. *Chaos, Solitons & Fractals*, 2008, **37**(3), 918–927.
- Carbone, A., Castelli, G. and Stanley, H.E., Time-dependent Hurst exponent in financial time series. *Physica A*, 2004, **344**(1/2), 267–271.

- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Christie-David, R., Chaudhry, M. and Koch, T.W., Do macroeconomics news releases affect gold and silver prices? *J. Econ. Bus.*, 2000, **52**, 405–421.
- Ciner, C., On the long run relationship between gold and silver prices. *Glob. Financ. J.*, 2001, **12**, 299–303.
- Clark, A., The use of Hurst and effective return in investing. *Quant. Finance*, 2005, **5**(1), 1–8.
- Cochran, S.J., Mansur, I. and Odusami, B., Volatility persistence in metal returns: A FIGARCH approach. *J. Econ. Bus.*, 2012, **64**(4), 287–305.
- Du, G. and Ning, X., Multifractal properties of Chinese stock market in Shanghai. *Physica A*, 2005, **387**(1), 261–269.
- Ellis, C. and Hudson, C., Scale-adjusted volatility and the Dow Jones index. *Physica A*, 2007, **378**(2), 374–384.
- Eom, C., Choi, S., Oh, G. and Jung, W.S., Hurst exponent and prediction based on weak-form efficient market hypothesis of stock markets. *Physica A*, 2008, **387**(18), 4630–4636.
- Escribano, A. and Granger, C.W.J., Investigating the relationship between gold and silver prices. *J. Forecast.*, 1998, **17**, 81–107.
- Fama, E.F., Market efficiency, long term returns, and behavioural finance. *J. Financ. Econ.*, 1998, **49**, 283–306.
- Figuerola-Ferretti, I. and Gonzalo, J., Price discovery and hedging properties of gold and silver markets. Working Paper presented at the 2010 European Financial Management Conference, Aarhus. Paper dated January 19, 2010.
- Grech, D. and Pamula, G., The local Hurst exponent of the financial time series in the vicinity of crashes on the Polish stock exchange market. *Physica A*, 2008, **387**(16/17), 4299–4308.
- Hurst, H.E., Long-term storage capacity of reservoirs. *Trans. Am. Soc. Civ. Engrs*, 1951, **116**, 770–799.
- Jiang, Z.Q. and Zhou, W.X., Multifractality in stock indexes: Fact or Fiction? *Physica A*, 2008, **387**(14), 3605–3614.
- Johansen, S., Cointegration and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 1991, **59**(6), 1551–1580.
- Lo, A.W., Long-term memory in stock market prices. *Econometrica*, 1991, **59**(5), 1279–1313.
- Lucey, B. and Tully, E., Seasonality, risk and return in daily COMEX gold and silver data 1982–2002. *Appl. Financ. Econ.*, 2006, **16**, 319–333.
- Mandelbrot, B.B., Stochastic volatility, power laws and long memory. *Quant. Finance*, 2001, **1**(6), 558–559.
- Mandelbrot, B.B. and Wallis, J.R., Robustness of the rescaled range R/S in the measurement of noncyclic long-run statistical dependence. *Water Resour. Res.*, 1969, **5**(5), 967–988.
- Matos, J., Gama, S., Ruskin, M.A., Sharkasi, H. and Crane, M., Time and scale Hurst exponent analysis for financial markets. *Physica A*, 2008, **387**(15), 3910–3915.
- McCauley, J.L., Bassler, K.E. and Gunaratne, G.H., Martingales, nonstationary increments, and the efficient market hypothesis. *Physica A*, 2008, **387**(15), 3916–3920.
- Mohanty, S., Nandha, M. and Bota, G., Oil shocks and stock returns: The case of the Central and Eastern European (CEE) oil and gas sectors. *Emerg. Mkts Rev.*, 2010, **11**(4), 358–372.
- Mulligan, R., Fractal analysis of highly volatile markets: An application to technology equities. *Q. Rev. Econ. Finance*, 2004, **44**, 155–179.
- Owczarczuk, M., Long memory in patterns of mobile phone usage. *Physica A*, 2012, **391**(4), 1428–1433.
- Szilagyi, P.G. and Batten, J.A., Covered interest parity arbitrage and temporal long-term dependence between the US dollar and the Yen. *Physica A*, 2007, **376**, 409–421.
- Taguchi, H., Monetary autonomy in emerging market economies: The role of foreign reserves. *Emerg. Mkts Rev.*, 2011, **12**(4), 371–388.

- Takami, Y., Tabak, B. and Miranda, M., Interest rate option pricing and volatility forecasting: An application to Brazil. *Chaos, Solitons & Fractals*, 2008, **38**(3), 755–763.
- Urich, T.J., Modes of fluctuation in metal futures prices. *J. Fut. Mkts*, 2000, **20**, 219–241.
- Wang, Y., Wei, Y. and Wu, C., Analysis of the efficiency and multifractality of gold markets based on multifractal detrended fluctuation analysis. *Physica A*, 2011a, **390**(5), 817–827.
- Wang, Y., Wei, Y. and Wu, C., Detrended fluctuation analysis on spot and futures markets of West Texas Intermediate crude oil. *Physica A*, 2011b, **390**(5), 864–875.

12

Gold and the U.S. Dollar: Tales from the Turmoil

Paolo Zagaglia

Massimiliano Marzo

12.1	Introduction	227
12.2	The Dataset	230
12.3	Some Structural Evidence on Volatility Spillovers.....	232
12.4	A Look at the Tails.....	236
12.5	Conclusion	240
	Acknowledgements	241
	References.....	241

We investigate how the relation between gold prices and the U.S. dollar has been affected by the recent turmoil in financial markets. We use spot prices of gold and spot bilateral exchange rates against the euro and the British pound to study the pattern of volatility spillovers. We estimate the bivariate structural GARCH models proposed by Spargoli and Zagaglia to gauge the causal relations between volatility changes in the two assets. We also apply the tests for change of codependence of Cappiello *et al.* to study the impact of the turmoil on the relation between gold and the U.S. dollar. We document the ability of gold to generate stable comovements with the dollar exchange rate that have survived the recent phases of market disruption. Our findings also show that exogenous increases in market uncertainty have tended to produce reactions of gold prices that are more stable than those of the U.S. dollar.

Keywords: Exchange rates; Gold; Multivariate GARCH; Quantile regressions

JEL Classification: C2, C22, F3, F31, F33

‘(Gold’s) quantity cannot increase at the same rate as you can print money, which will eventually weaken the US dollar’, Faber said on Thursday in a live interview.

Marc Faber on CNBC, March 4, 2010

‘There is still a link on a day-to-day and weekly basis between gold and the dollar, so if the euro weakens gold will come under pressure’, said Standard Bank analyst Walter de Wet. ‘We might test down towards 1.10\$ if the euro goes to 1.30\$’.

CNBC Commentary, March 4, 2010

12.1 Introduction

Various accounts of the recent financial turmoil have stressed the propensity of investors to turn away from risky securities into ‘safe’ assets. This flight to safety has also taken the form of a renowned interest in gold as an asset class. In fact, it is typically argued that the price of this precious metal is uncorrelated

with both stock and bond prices during episodes of market crash (see, e.g. Baur and Lucey (2010) and Baur and McDermott (2010)). Gold is also identified as a 'hedge' against fluctuations in the U.S. dollar (USD) on average (Baur and Lucey 2010).^{*} The properties of safe asset and hedging capabilities suggest that the dollar price of gold should increase when the bilateral exchange of the U.S. dollar against other currencies depreciates.

The international role of gold dates back to the early years of the nineteenth century when several countries adopted the 'gold standard'. In this arrangement, the value of a currency was backed by government holdings of gold. The gold standard collapsed in 1971 with the decision of President Richard Nixon to end the dollar convertibility into gold.

Cappie *et al.* (2005) point out that two reasons are typically suggested for the use of gold as a hedging instrument or safe asset against exchange rate risk. First, a number of financial products are available that track the price of gold, despite the fact that they do not involve the property of the physical commodity. For instance, there are commodity Exchange Traded Funds that are linked to gold. Second, gold is often pointed at as a protection against currency fluctuations worldwide, not only for the U.S. dollar (see also Sjaastad and Scacciavillani (1996)). The results of Cappie *et al.* (2005) show that the hedging power of gold for the U.S. dollar has varied widely since 1971. In particular, they argue that the degree of protection offered by the dollar depends on largely unpredictable events.[†]

In this paper, we study how the relation between gold prices and the U.S. dollar has been affected by the turmoil that erupted in financial markets in 2007. In other words, we provide evidence on whether gold has behaved as a safe asset against currency fluctuations since 2007.

The tests for gold as a safe asset or as a hedging instrument consider the patterns of conditional correlation between gold returns and the returns of alternative assets, and how these change through time. A negative correlation is indicative of hedging capabilities. We consider a more general approach, and shed light on properties of the relation between gold and the U.S. dollar that have not been considered previously. We provide a new definition based on the evolution of the pattern of 'contagion', on one hand, and 'comovements', on the other, during the turmoil period. This allows us to reinterpret some previous results available in the literature under a novel perspective, using methods that shed new light on old issues.

The idea of contagion is widely used in empirical studies of financial market crises. The aim is to investigate how exogenous changes in volatility are transmitted between markets. We believe that this is an issue of major importance for understanding both the sources and the channels through which risk propagates. Gauging how volatility 'spills over' across assets requires making structural assumptions about the relations between two assets.[‡] The validity of these assumptions can then be verified both through standard statistical tests, as well as by considering whether their model implications are reasonable.

Our structural approach differs noticeably from the 'reduced-form' evidence on the hypothesis of hedging tool or safe asset for gold. An example of an application for this method can be found, amongst others, in Cappie *et al.* (2005). This key study uses information from the statistical coefficient of univariate generalized autoregressive conditional-heteroskedasticity (GARCH) models to come to conclusions on the hedging properties of gold for the U.S. dollar. In our paper, we dive into the question of how large the volatility spillovers are that arise from changes to gold returns onto the USD. Our method allows us to study the role of specific channels of transmission of instability across asset markets.

^{*} Several contributions indicate that gold provides stability to industry portfolios. For instance, Davidson *et al.* (2003) show that standard international asset pricing models prescribe a systematic exposure to gold.

[†] We should stress that the price of gold is also affected by some of the driving factors of the dollar exchange rate, such as changes in the inflation rate (see, e.g. McCown and Zimmerman (2006)) and the release of macroeconomic news (Christie-David *et al.* (2000)).

[‡] By 'structural' we refer to the hypothesis on the distinction between the causes of purely exogenous changes of variables and the endogenous impact of these shocks.

Since our sample data includes the recent financial turmoil, we also consider the joint tail behaviour of both gold and the USD returns. To put it differently, we investigate how the relation between gold and the USD is affected by extreme events within an episode of financial-market turbulence.

To deal with these two main questions, this paper employs two different econometric frameworks to evaluate the change in the comovements between gold and the U.S. dollar. First, we provide evidence on how the impact of volatility shocks and spillovers has changed during the turmoil. We estimate the standard bivariate GARCH models proposed by Engle and Kroner (1995). In order to uncover the role played by market linkages in the propagation of volatility shocks, we consider the extension to the structural BEKK model discussed by Spargoli and Zagaglia (2008). This amounts to using the information from time-varying conditional heteroskedasticity to identify causal relations between the volatility movements in the price of gold and the U.S. dollar (Rigobon 2003).

In the following step, we study the evolution of comovements between extreme prices. We use the measure of contagion proposed by (Cappiello *et al.*, 2005) to investigate whether the probability of observing closer comovements has increased since August 2007. The framework of Cappiello *et al.* (2005) is based on the computation of the probability of a variable falling below a threshold conditional on the same pattern for the other variable. Thresholds are obtained through quantile estimation. In this statistical model, a high conditional probability of comovement implies a strong codependence between the variables.

Our results indicate that the outbreak of the turmoil in August 2007 has not increased the uncertainty of either gold, or the USD. Rather, the triggering event for market volatility was the bankruptcy of Lehman Brothers in September 15, 2008. The volatility links between gold and the foreign exchange market exacerbate the transmission of shocks and the resulting uncertainty in price movements.

We uncover two features of gold that are disregarded in the literature. We show that exogenous increases in market uncertainty have tended to produce reactions of gold prices that are more stable than those of the U.S. dollar during the turmoil period. Hence, gold prices are affected by market uncertainty to a smaller extent than the USD exchange rate. For a given level of correlation between these two assets, this suggests that the volatility shock originating from gold prices and transmitted to the USD has a more limited size. Overall, this is a source of value added for gold in currency portfolios.

Furthermore, we uncover a feature of gold that is typically disregarded. This consists of its ability to generate comovements with the U.S. dollar exchange rate that are stable through time. Strikingly, such stability has survived the extreme events that have taken place during the recent phases of market disruption. Also, it characterizes both tail episodes of abrupt market swings, as well as conditions of milder changes. In other words, even though our empirical methods focus on a ‘crisis within a crisis’, we confirm the previous findings suggesting that gold can be considered a safe asset.

The results presented here have important implications for risk management. In this area, the use of the concept of volatility shocks is both appealing from a methodological point of view, and for the purpose of practical applications. Understanding the determinants of conditional volatility is relevant for modelling the risk factors and how these manifest themselves. For instance, in the context of risk budgeting for currency portfolios, we should expect gold to provide a contribution to portfolio risk—following an increase in market-wide uncertainty—smaller than that generated by the USD. There are also implications in terms of portfolio management strategies. Including holdings of gold and selected USD-bilateral rates removes a source of exchange-rate risk, namely the risk of transmission of volatility shocks. In particular, we suggest that gold shields a portfolio from the volatility spillovers potentially arising from movements in the U.S. dollar. This feature does not depend on the severity of market fluctuations.

The paper is organized as follows. [Section 12.2](#) presents the dataset and discusses some properties of the series. [Section 12.3](#) elaborates a structural BEKK model to provide an interpretation of the effects of volatility changes. In [Section 12.4](#) we focus on the relationship at the tails of the distributions, and present evidence of comovement patterns of extreme price changes based on quantile regressions. [Section 12.5](#) concludes the paper.

12.2 The Dataset

In this paper we use daily data for spot contracts of exchange rates between the U.S. dollar and the euro and the U.S. dollar and the British pound. Our dataset also includes prices for spot contracts of gold negotiated at the Chicago Board of Trade, expressed in U.S. dollars for 100 ounces.* Both the exchange rate and gold price data are extracted from Bloomberg. The dataset spans from October 13, 2004 to March 5, 2010.† We consider August 9, 2007 as the starting date for the outbreak of the turmoil in financial markets worldwide. During that day, BNP Paribas froze the redemption of three investment funds, and the resulting panic forced the European Central Bank to start extraordinary measures for the supply of liquidity in the euro interbank market.

Figure 12.1 plots the data series (in logarithms), and Table 12.1 reports some descriptive statistics. During the turmoil, the U.S. dollar depreciates, on average, with respect to both the euro and the British pound. The average return of gold increases as well. The returns on both exchange rates and gold are




FIGURE 12.1 Data series.

* In order to deal with the problematic issue of asynchronous trading, we use both gold prices and exchange rates sampled at 1.30 p.m. ET. This time identifies the closing of the open outcry session for gold derivatives on weekdays.

† We choose October 13, 2004 as the starting observation in our dataset because we intend to use a roughly equal amount of time-series information to characterize the periods before and during the turmoil experience. In other words, this is only meant to avoid the possible distortions that may arise if the sample is skewed towards tranquil periods for financial markets.

more volatile during the turmoil. Interestingly, the increase in volatility for the exchange rates is far higher than for gold. The changes in the kurtosis coefficients confirm that the USD exchange rate is subject to events more extreme than gold.

It is also instructive to consider the two scatter plots of the (logs) bilateral exchange rates versus the price of gold. Figure 12.2 shows the plots, emphasizing the periods before and during the turmoil. The

TABLE 12.1 Descriptive Statistics of the Asset Returns (In Per cent)

	Before the turmoil				During the turmoil			
	Mean	St. dev.	Skewness	Kurtosis	Mean	St. dev.	Skewness	Kurtosis
Gold	0.0687	1.1287	-0.8199	7.6902	0.0761	1.6003	0.1552	5.4622
USD/EUR	0.0165	0.4834	-0.0014	3.6719	-0.0019	0.8002	0.3979	6.3870
USD/GBP	0.0178	0.4816	0.0819	3.5273	-0.0420	0.8942	-0.3278	7.5810




FIGURE 12.2 Scatter plots.

USD/EUR exhibits a positive correlation with the price of gold both before and during the turmoil. This can be interpreted as evidence suggesting that investors have used gold as a hedge for the USD/EUR exchange rate. For the USD/GBP exchange rate though, the turmoil appears to be characterized by a break in the correlation with the price of gold. The turmoil has apparently induced investors to hold gold for hedging purposes also against a fall of the USD/GBP exchange rate.

12.3 Some Structural Evidence on Volatility Spillovers

In order to understand how the hedging power of gold changes during the turmoil, we investigate how the pattern of transmission of shocks with the U.S. dollar evolves. For that purpose, we estimate the structural BEKK model discussed by Spargoli and Zagaglia (2008).

We assume that the joint evolution of a bilateral USD exchange rate return and gold returns can be summarized by a structural vector-autoregressive (VAR) model,

$$Ax_t = \Psi + \Phi(L)x_t + \eta_t, \quad (12.1)$$

where x_t is the vector of the vector of returns, ψ is a vector of constants, A is a matrix of structural parameters, and $\eta_t \sim N(0, h_t)$ is a vector of structural shocks. The structural innovations exhibit conditional heteroskedasticity. We use the BEKK-GARCH model of Engle and Kroner (1995),

$$h_t = CC' + Gh_{t-1}G' + T\eta_{t-1}\eta_{t-1}'T'. \quad (12.2)$$

In model (12.1)–(12.2), the regressors are not exogenous because their source of variation is represented by the dependent variable in the same equation through another equation in the system. In order to achieve identification of the relations modelled in the VAR we rely on heteroskedasticity. This idea was originally introduced by Wright (1928) and recently developed by Rigobon (2003). The heteroskedasticity approach to identification amounts to using the information from time-varying volatility as a source of exogenous variation in the endogenous variables. To see this, let us consider the reduced-form VAR model

$$x_t = c + F(L)x_t + v_t, \quad (12.3)$$

where $c = A^{-1}\Psi$, $F(L) = A^{-1}\Phi(L)$ and $v_t = A^{-1}\eta_t$ are the reduced-form innovations, whose variance–covariance matrix is a combination of the variance–covariance matrix of the structural-form innovations, that is

$$H_t = Bh_tB', \quad (12.4)$$

$$H_t = BCC'B' + BGh_{t-1}G'B' + BT\eta_{t-1}\eta_{t-1}'T'B'. \quad (12.5)$$

In this formulation the variance–covariance matrix of the reduced-form innovations is a function of the structural innovations, which the econometrician does not know. However, we can use the equality to show that

$$\eta\eta' = Av_t v_t' A', \quad (12.6)$$

and represent it in terms of the reduced-form innovations as

$$H_t = BCC'B' + BGAH_{t-1}A'G'B' + BTA v_{t-1} v_{t-1}' A'T'B'. \quad (12.7)$$

This reduced form is then used for the estimation.

After estimating the model, we compute impulse-response functions. In structural GARCH models, these functions show the impact that a shock produces on the conditional second moments of the variables in the system. However, differently from the impulse-response functions for a standard VAR, the impulse responses of a structural GARCH depend both on the magnitude of the shock and on the period during which the shock itself takes place. This is due to the fact that the residuals enter the model in quadratic form. Hence, differently from the case of linear models, the magnitude of the effects of a shock is not proportional to the size of the shock itself. This allows us to compute a distribution of impulse responses following each shock. For this purpose, we use the concept of Volatility Impulse Response Functions (VIRF) proposed by Hafner and Herwartz (2006). The impulse-response function for a vech-GARCH model can be written as

$$V_t(\xi_0) = E(\text{vech}(H_t) | \xi_0, I_{t-1}) - E(\text{vech}(H_t) | I_t). \quad (12.8)$$

The response at time t of the variances and covariances following a shock η_t in $t = 0$, denoted $V_t(\xi_0)$, is equal to the difference, conditioned on the information set I_{t-1} at time $t-1$ and on the shock η_0 , of the variance (or covariance) at t from its expected value conditional on the information set of the previous period.

We used standard likelihood methods to estimate two structural BEKK(p, q) models.* The variables modelled are the returns of each bilateral exchange rate and the return on gold. Each model has order $p=1$ and $q=1$.† The parameter estimates are listed in Table 12.2. The standard errors are computed using the delta method as in Spargoli and Zagaglia (2008). It should be stressed that the estimated parameters for the BEKK are statistically significant at standard confidence levels.‡

Figure 12.3 reports the conditional variances of both the structural BEKK model, as well as the variances from a reduced-form model that disregards the issue of identification of causal relations. Two observations emerge. The first is that the outbreak of the turmoil alone did not enhance the pattern of

TABLE 12.2 Parameter Estimates of the Structural BEKK.

Parameter	USD/EUR–gold		USD/GBP–gold	
	Point estimate	t-Stat.	Point estimate	t-Stat.
$c_{1,1}$	0.2908	3.6784	0.2908	2.7094
$c_{1,2}$	1.0144	12.8313	0.3608	1.5084
$c_{2,2}$	2.1310	26.9553	2.6808	3.5785
$a_{1,2}$	−4.2570	−53.8473	2.8468	2.0312
$a_{2,1}$	2.8960	36.6318	1.8961	7.6318
$g_{1,1}$	0.8438	10.6733	−0.7600	9.0153
$g_{1,2}$	−0.2194	−2.7752	0.7600	2.4137
$g_{2,1}$	−0.1291	−1.6330	0.7600	6.2190
$g_{2,2}$	0.7600	9.6133	−0.7600	4.6201
$t_{1,1}$	0.4715	5.9641	0.6100	3.4309
$t_{1,2}$	0.2745	3.4722	0.7745	6.0121
$t_{2,1}$	0.2399	3.0345	0.2935	4.0922
$t_{2,2}$	0.3840	4.8573	0.5945	4.8022

* We maximize the likelihood functions by simulated annealing.

† We select the model that delivers the largest value for the likelihood.

‡ Since we use a derivative-free optimization method for the likelihood, we compute standard errors through the ‘delta’ method.




FIGURE 12.3 Estimated reduced-form and structural variances.

uncertainty characterizing the returns of either gold or the bilateral exchange rates. Rather, there is a surge in volatility in correspondence with the bankruptcy of Lehman Brothers of September 15, 2008. The second point is that the reduced-form estimates lay on top of the structural variances. This suggests that the links between gold and the foreign exchange markets exacerbate the transmission of shocks and, thus, the uncertainty in the price movements.

Figure 12.4 plots the structural and reduced-form estimates of the correlations. The turmoil has no clear-cut effect on the dynamics of the correlations. On the other hand, following the Lehman crash, there is an increase in the average reduced-form correlations, in comparison with the structural correlations, that lasts until the beginning of 2009. This indicates that the linkages between markets enhance the tendency of the assets to move together.

In order to gain some understanding of the transmission of shocks, Figure 12.5 reports the means of the distributions of volatility-impulse responses following a one standard deviation shock. The returns of gold are stable, in the sense that the marginal change in their uncertainty is not substantially affected by the turmoil. The reaction of the variance of the bilateral exchange rates is, instead, magnified.




FIGURE 12.4 Estimated reduced-form and structural correlations.

Figure 12.5 also shows that the price of gold and the bilateral exchange rates tend to decouple as the response of the covariance on impact falls for the turmoil period.

There are two issues with this type of introductory evidence. We are looking at the mean of the distribution of volatility-impulse responses. There may be relevant information that is carried over in the relation between gold and foreign exchange rates at the tail of the distribution. The second aspect concerns the lack of a criterion of statistical significance of the figures discussed earlier. Both these issues are addressed in the next section, where we consider a framework for testing how the tail relation between gold and foreign exchange rates has changed after the turmoil.

12.4 A Look at the Tails

Standard tests for comovements rely on the estimation of correlations between asset returns. These tests are, however, typically significant both to the presence of heteroskedasticity, and to departures from normality in the empirical distributions of two returns. The comovement box of Cappiello *et al.* (2005) relies on semiparametric methods to provide a robust method for analysing comovements.




FIGURE 12.5 Mean of the distribution of volatility-impulse responses before and after August 9, 2007.

Let $\{r_{i,t}\}_{t=1}^T$ and $\{r_{j,t}\}_{t=1}^T$ denote the time series of returns on two different assets, namely a USD exchange rate and gold. Define by $q_{\theta,i}^r$ the θ -quantile of the conditional distribution of $r_{i,t}$ at time t . $F_t(r_i, r_j)$ denotes the conditional cumulative joint distribution of the two asset returns. Finally,

$$F_{it}^-(r_i | r_j) := \text{prob}(r_{i,t} \leq r_i | r_{j,t} \leq r_j), \quad (12.9)$$

$$F_{it}^+(r_i | r_j) := \text{prob}(r_{i,t} \geq r_i | r_{j,t} \geq r_j), \quad (12.10)$$

The conditional probability,

$$p_t(\theta) := \begin{cases} F_{it}^-(q_{\theta,t}^r | q_{\theta,t}^r) & \text{if } \theta \leq 0.5, \\ F_{it}^+(q_{\theta,t}^r | q_{\theta,t}^r) & \text{if } \theta > 0.5, \end{cases} \quad (12.11)$$

can be used to represent the characteristics of $F_t(r_i, r_j)$. In fact, $p_t(\theta)$ measures the probability that the returns at maturity i are below its θ -quantile, conditional on the same event occurring at maturity j .

The information about $p_t(\theta)$ is summarized in the so-called ‘comovement box’. This is a square with unit size where $p_t(\theta)$ is plotted against θ . Since the shape of $p_t(\theta)$ depends on the joint distribution of the two time series, it can be derived only by numerical simulation.*

The framework of Cappiello *et al.* (2005) can also be used to test whether the dependence between two markets has changed over time. Given a cut-off date of a specific event, we can estimate the conditional probability of comovements in two different periods, and plot the estimated probabilities in a graph. Differences in the intensity of comovements can then be detected. This idea can be formalized in a simple way. Denote by $p^A(\theta) := A^{-1} \sum_{t < \tau} p_t(\theta)$ and $p^B(\theta) := B^{-1} \sum_{t \geq \tau} p_t(\theta)$ the average conditional probabilities before and after a certain event occurs at a threshold τ , with A and B the number of corresponding observations. Let $\Delta(\underline{\theta}, \bar{\theta})$ denote the area between $p^A(\theta)$ and $p^B(\theta)$. A measure of ‘contagion’ or ‘spillovers’ between the two markets can be obtained by noting that contagion increases if

$$\Delta(\underline{\theta}, \bar{\theta}) = \int_{\underline{\theta}}^{\bar{\theta}} [p^B(\theta) - p^A(\theta)] d\theta > 0. \quad (12.12)$$

We stress that, unlike the standard measures of correlation, $\Delta(\underline{\theta}, \bar{\theta})$ allows us to study changes in codependence over specific quantiles of the distribution.

Several steps are followed to construct the comovement box and test for differences in conditional probabilities. First, we estimate univariate time-varying quantiles using the Conditional Autoregressive Value at Risk (CAViaR) model proposed by Engle and Manganelli (2004). For each series and each quantile, we create an indicator variable that takes the value one if the return is lower than this quantile, and zero otherwise. Then we regress the θ -quantile indicator variable on market j on the θ -quantile indicator on market i . The estimated regression coefficients provide a measure of the conditional probabilities of comovements, and of their changes across regimes.

The time-varying quantiles of the returns are estimated using the CAViaR model of Engle and Manganelli (2004). The quantiles of the returns r_t are assumed to follow the autoregressive model

$$q_t(\beta_\theta) = \beta_{\theta,0} + \sum_{i=1}^q \beta_{\theta,i} q_{t-i} + \sum_{i=1}^p l(\beta_{\theta,j}, r_{t-j}, \Omega_t), \quad (12.13)$$

* The interested reader can refer to Cappiello *et al.* (2005, p. 9) for a thorough discussion on the comovement box.

where Ω_t denotes the information set at time t . The autoregressive terms of the quantiles are meant to capture the clustering of volatility that is typical of financial variables. Including a predetermined information set allows us instead to consider the interaction between the quantiles and the conditions of the market. Following Cappiello *et al.* (2005), we estimate the time-varying quantiles using the following specification of the CAViaR:

$$q_t(\beta_\theta) = \beta_{\theta,0} + \beta_{\theta,1}d_t + \beta_{\theta,2}r_{t-1} + \beta_{\theta,3}q_{t-1}(\beta_\theta) - \beta_{\theta,2}\beta_{\theta,3}r_{t-2} + \beta_{\theta,4}|r_{t-1}|. \quad (12.14)$$

The dummy variable d_t ensures that the periods of high and low volatility have the same proportion of quantile exceedances.

In order to investigate the specification of the CAViaR model, we compute the DQ test of Engle and Manganelli (2004). The null hypothesis of the DQ test consists of the lack of autocorrelation in the exceedances of the quantiles. Figure 12.6 reports the p -values for 99 conditional quantiles, together with the p -values for unconditional quantiles. The specification with unconditional quantiles is rejected over the entire domain.

Figure 12.7 plots the estimates of the conditional probabilities of comovements before and during the turmoil period. We also report confidence bands of plus/minus twice the standard errors around the




FIGURE 12.6 p -Values of the dynamic quantile test.




FIGURE 12.7 Estimated tail codependence.

TABLE 12.3 Test of Difference in Tail Coincidences between Periods Before and During the Turmoil

	Lower tail: $\theta \leq 0.5$		Higher tail: $\theta \geq 0.5$	
	$\hat{\delta}(0, 0.5)$		$\hat{\delta}(0.5, 1)$	
	Stat.	s.e.	Stat.	s.e.
USD/EUR–gold	1.7555	2.1640	0.4084	2.5282
USD/GBP–gold	2.2261	2.7650	1.0981	2.3332

estimates of the probabilities of comovements. Two observations emerge from [Figure 12.7](#). The first is that extreme events due to the turmoil provide relevant information on the relation between the USD exchange rate and gold prices. This can be seen by comparing the changes in the positions of the tails of the distributions before and during the turmoil. In other words, the relation of comovements between the USD exchange rates and gold changes across time, especially when extreme events materialize in financial markets. The second point of interest is that the probability of comovement induced by the turmoil is not statistically different from the comovement probabilities characterizing the pre-turmoil period. In more intuitive terms, this suggests that the transmission mechanism of shocks between bilateral exchange rates and the price of gold is not affected by the turmoil.

Table 12.3 reports the results of the quantile tests for contagion for two parts of the distributions. All the test statistics have a positive sign, indicating an increase in comovement after the turmoil, on average. However, the null hypothesis of conditional comovement between the USD exchange rates and gold is rejected strongly in all cases. These findings uncover an important source of stability generated by gold. In this paper we stress the ability of this asset to shield portfolios from the effects of spillovers due to swings in the U.S. dollar, independently from how extreme the market conditions are.

12.5 Conclusion

The relation between gold prices and the U.S. dollar exchange rate has been the subject to intense scrutiny. Using recently developed econometrics methods, we reconsider the case for gold as a safe asset against fluctuations in the USD. Following the account of numerous commentators, we investigate whether the beneficial properties of holding gold have survived the recent turbulence in financial markets that started in August 2007.

Using methods from the literature on financial crises, we study the impact of the turmoil on two sources of these properties. We consider the evolution of the pattern of volatility spillovers between gold prices and the dollar. We investigate whether the turmoil has caused an increase in contagion, defined as the probability of extreme tail events in both gold prices and the dollar.

On general grounds, our results confirm the widely held belief that gold can be considered a financial asset. Several qualifications apply though. We find that the defining moment in the transmission of volatility shocks between gold and the USD coincides with the Lehman Brothers bankruptcy. Even if we consider the tails of the bivariate distributions, we find that holding gold generates patterns of comovements with the dollar that have survived the recent episodes of market turbulence. Most of all, we show that exogenous volatility shocks tend to generate reactions of gold prices that are more stable than those of the U.S. dollar. In other words, gold prices react to market uncertainty to a smaller extent than the USD. Since a smaller reaction of volatility is then transmitted to the USD, this pins down a key source of stability for gold holdings in currency portfolios.

To our knowledge, this is the first contribution that discusses the structural features of volatility transmission in the gold market. The relevant question concerns the nature of the determinants of the changes in contagion. Thus, the analysis of this paper can be extended along different and, yet, promising avenues for future research. For instance, it would be important to understand how deep the relation between gold prices and the U.S. dollar is around the business cycle. We can also consider the issue of whether the properties of gold discussed earlier arise also in relation to other asset classes. We might expect U.S. stocks to represent a relevant candidate.

The available literature has not yet quantified the advantages of gold holdings in a systematic framework of portfolio allocation when extreme events take place. In this sense, it would be interesting to compute the implications for standard risk measures, such as the Value-at-Risk or the expected tail loss. Other types of multivariate models of tail behaviour could be used. In a related work-in-progress, we consider bivariate copulas to model the joint determinants of co-exceedances. Finally, the recent experience of financial market turbulence suggests that structural breaks in the conditional mean, variance and correlations are important factors that should be accounted for in the relation between assets. This hints at the advantages of using models with time varying parameters for understanding the portfolio implications of changing relations between gold prices and the U.S. dollar.

Acknowledgements

We are very grateful to two anonymous referees for their helpful comments and suggestions on a previous draft. This research was carried out by the authors for purely academic purposes and without financial support from any private institution.

References

- Baur, D.G. and Lucey, B.M., Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. *Financ. Rev.*, 2010, 5(3), 217–229.
- Baur, D.G. and McDermott, T.K., Is gold a safe haven? International evidence. *J. Bank. Finance*, 2010, forthcoming.
- Cappie, F., Mills, T.C. and Wood, G., Gold as a hedge against the dollar. *J. Int. Financ. Mkts, Inst. Money*, 2005, 15, 343–352.
- Cappiello, L., Gerard, B. and Manganelli, S., Measuring comovements by regression quantiles. ECB Working Paper 501, 2005.
- Christie-David, R., Chaudhry, M. and Koch, T.W., Do macroeconomics news releases affect gold and silver prices. *J. Econ. Bus.*, 2000, 52(5), 405–421.
- Davidson, S., Faff, R. and Hillier, D., Gold factor exposures in international asset pricing. *J. Int. Financ. Mkts, Inst. Money*, 2003, 13, 271–289.
- Engle, R.F. and Kroner, F.K., Multivariate simultaneous generalized ARCH. *Econometr. Theory*, 1995, 11, 122–150.
- Engle, R. and Manganelli, S., CAViaR: Conditional autoregressive value at risk by regression quantiles. *J. Bus. Econ. Statist.*, 2004, 22, 367–381.
- Forbes, K.J. and Rigobon, R., No contagion, only interdependence: Measuring stock market comovements. *J. Finance*, 2002, 57, 2223–2261.
- Hafner, C.M. and Herwartz, H., Volatility impulse responses for multivariate GARCH models: An exchange rate illustration. *J. Int. Money Finance*, 2006, 25(5), 719–740.
- McCown, J.R. and Zimmerman, J.R., Is gold a zero-beta asset? Analysis of the investment potential of precious metals. Working Paper Series, Oklahoma City University, 2006.
- Rigobon, R., Identification through heteroskedasticity. *Rev. Econ. Statist.*, 2003, 85(4), 777–792.
- Rigobon, R. and Sack, B., Measuring the reaction of monetary policy to the stock market. *Q. J. Econ.*, 2003, 118(2), 639–669.
- Sjaastad, L.A. and Scacciavillani, F., The price of gold and the exchange rate. *J. Int. Money Finance*, 1996, 15(6), 879–897.
- Spargoli, F. and Zagaglia, P., The comovements along the forward curve of natural gas futures: A structural view. *J. Energy Mkts*, 2008, 1(3), 23–47.
- Wright, P., *The Tariff on Animal and Vegetable Oils*, 1928 (Macmillan: New York).

13

A Flexible Model of Term-Structure Dynamics of Commodity Prices: A Comparative Analysis with a Two-Factor Gaussian Model

13.1	Introduction	244
13.2	Comparison of Two Modelling Approaches for the Term Structure of Commodity Futures.....	245
	A Flexible Model of Commodity Futures Returns • A Conventional Two- Factor Term-Structure Model of Commodity Prices • Model Comparison	
13.3	Data and Estimation.....	250
	Data • Estimation Results: Model Specification • Estimation Results: Flexible Model • Estimation Results: Conventional Two-Factor Model • Estimation Results: Composite Model	
13.4	Implications for the Optimal Hedging Strategy.....	270
	Optimal Hedging • Optimal Hedging Strategy According to the Three Models	
13.5	Conclusion	276
	Acknowledgements	277
	References.....	277

Hiroaki Suenaga

This study compares two approaches to modelling a term structure of commodity prices. The first approach specifies the stochastic process of the underlying spot price and derives from the stipulated spot price dynamics valuation formulas of futures and other derivative contracts through no arbitrage. The second approach, as introduced by Smith [*J. Appl. Econometr.*, 2005, **20**, 405–422], is to model the dynamics of the entire futures curve directly by a set of common stochastic factors and to specify factor loadings by flexible functions of time-to-maturity and contract delivery month. Empirical applications of the models to four commodities (gold, crude oil, natural gas and corn) reveal that the volatility of futures prices exhibits more complex dynamics than the pattern implied by the model stipulating a two-factor Gaussian process of the underlying spot price. Specifically, the flexible model of futures returns depicts the maturity effect and, particularly for the three consumption commodities, strong seasonal and cross-sectional variations in variance

and covariance of concurrently traded contracts. Incorporating the depicted variance and covariance dynamics leads the flexible model of futures returns to suggest hedging strategies that are more effective than the strategies based on the conventional two-factor Gaussian model.

Keywords: Commodity prices; Term structure; Volatility modelling; Hedging techniques

JEL Classification: C2, C23, G1, G19

13.1 Introduction

Recent increase in the level and volatility of oil, metals and other primary commodity prices has created tremendous uncertainties for producers, consumers and other traders of these commodities. Volatility of commodity prices also affects the national economy, both directly by altering revenue and expenditure on these commodities and indirectly by deferring new investments. In such circumstances, a better understanding of the stochastic properties of these commodity prices and tools to hedge against price risks becomes increasingly important for the smooth functioning of the commodity supply chain.

Stochastic dynamics of commodity prices and valuation of derivative contracts have long been studied in the field of financial economics. The standard approach in this literature is to specify the stochastic process of the underlying asset, usually the spot price of the commodity under investigation, and derive from the stipulated process valuation formulas of futures and other derivative contracts whose payoff depends on the value of the underlying asset realized at the contract maturity date (Hull 2002). This approach dates back to Black and Scholes (1973), who derived pricing formulas of European options under the assumption that the underlying asset value (stock price) follows a geometric Brownian motion (GBM). While following the same approach, many studies modelling commodity price dynamics commonly specify one of the underlying stochastic factors to follow a mean-reverting (MR) process because, for many consumption commodities, demand and/or supply response forces unusually high or low prices to revert to the long-run equilibrium level (Schwartz 1997).

Recent advances in this modelling approach have been attained through increasing the number of common stochastic factors and/or stipulating an increasingly complex stochastic process of each latent factor.* These flexible models generally exhibit a better fit to the observed price data while maintaining the model parsimonious with pricing formulas of derivative contracts typically determined by a small number of parameters characterizing the stochastic dynamics of the underlying factors. However, it is often understated that the benefit of a parsimonious specification is gained at the cost of potentially large errors in approximating true stochastic dynamics of commodity prices. It has been widely acknowledged that, unlike stocks and other conventional financial assets, commodity prices exhibit complex dynamics. The theory of storage (e.g. Williams and Wright (1991) and Routledge *et al.* (2000)) illustrates that, for a commodity with a significant storage cost, inter-temporal arbitrage establishes an equilibrium constellation of spot and futures prices along which the marginal benefit of current consumption is equal to or above the expected marginal benefit of storing a commodity for future consumption. The weak inequality stems from a non-negativity of physical storage. If supply is ample relative to demand, inter-temporal arbitrage induces positive inventory up to a point where the two prices differ by the cost of carry. In this case, the convenience yield, representing the implicit revenue from holding a physical asset, is close to zero. In contrast, when supply is scarce, discretionary inventory is driven to zero and the marginal benefit of current consumption exceeds the marginal benefit of future consumption due to high convenience yield. In this case, speculative storage plays a minor role in price determination and the inter-temporal price linkage breaks, which means that price correlations across concurrently

* Lautier (2005) provides a comprehensive review on applications of term-structure models to various commodities.

traded contracts vary by season. Many commodities also exhibit pronounced seasonality in price and volatility, reflecting seasonality in the underlying demand and/or supply. Volatility tends to be high in the period of tight demand-supply balance because market shocks of even a small magnitude can cause a large price swing. It is also expected that volatility is inversely related to inventory because demand and supply shocks can be absorbed through adjusting inventory. Stochastic processes of the spot price and other underlying factors stipulated in many models of commodity price dynamics, even recently developed complex models, are often too simple to induce a futures price formula that replicates the complex dynamics of commodity futures prices implied by the theory of storage.

An alternative approach to modelling a term structure of commodity prices, as recently introduced by Smith (2005) and later extended by Suenaga and Smith (2011), is to model directly the dynamics of futures curve. In this model, daily futures returns is decomposed into a set of common stochastic factors affecting all futures returns and an idiosyncratic term. By modelling futures *returns* rather than a price *level*, the model does not specify seasonal or other deterministic variation in the underlying spot price. This model also avoids specifying stochastic dynamics of the underlying factors and imposes no *a priori* restriction on the factor loadings that connect underlying factors to observed futures returns. Rather, the model specifies factor loadings and the variance of the idiosyncratic term directly by flexible functions so that they can replicate highly non-linear price dynamics of commodities with significant storage costs and seasonality in demand or supply.

In this study, I compare the two approaches to the modelling of the term structure of commodity prices; one specifying directly the dynamics of daily futures returns as a flexible function of common stochastic factors, and the other specifying the stochastic process of the underlying spot price. I apply the two models to futures price data from four commodity markets (crude oil, natural gas, gold and corn). Results from this empirical analysis illustrate that the volatility of daily futures prices exhibits highly non-linear dynamics that cannot be induced by the stochastic process of the underlying spot price stipulated in the conventional two-factor term-structure model. Specifically, the flexible model of futures returns depicts a maturity effect and, particularly for the three consumption commodities, strong seasonality in both its levels and compositions among the two common stochastic factors and the idiosyncratic error. These features together create substantial seasonal and cross-sectional variation in the price correlations of concurrently traded contracts. Incorporating the depicted dynamics of price volatility and cross-contract correlations allows the flexible model of futures returns to suggest hedging strategies that are more effective than a strategy based on the conventional term-structure model specifying a two-factor Gaussian process for the underlying spot price.

The next section presents the two approaches to modelling a term structure of commodity prices. The section also presents a composite model in which the factor loadings are specified as in the conventional spot-price-based approach, yet allows for a flexible variance structure of the idiosyncratic error. [Section 13.3](#) reports results from estimating the three models empirically with unbalanced panel data from four commodity markets. The results are compared across the models with an emphasis on seasonal and cross-sectional variations in the depicted price variance and cross-contract correlations. [Section 13.4](#) considers the implication of the models for an optimal strategy to hedge price risk. [Section 13.5](#) provides a synopsis of the findings to conclude the paper.

13.2 Comparison of Two Modelling Approaches for the Term Structure of Commodity Futures

This section presents a flexible model of futures curve dynamics in which daily futures returns is decomposed into a set of common latent factors and an idiosyncratic term. This model is then compared with the conventional term-structure model specifying stochastic dynamics of the underlying spot price and deriving pricing formulas of futures and other derivative contracts.

13.2.1 A Flexible Model of Commodity Futures Returns

One approach to modelling commodity price dynamics, as introduced by Smith (2005) and later extended by Suenaga and Smith (2011), is to model directly the daily price changes of all concurrently traded futures contracts.* In this approach, a daily return of futures contract is decomposed into the common latent factors and an idiosyncratic term. The model incorporates time-varying conditional heteroskedasticity of latent factors and time and cross-sectional variation in the factor loadings and idiosyncratic variances.

The model with two common factors is expressed in the following form:

$$\Delta \ln F_{m,t} = \theta_1(m, d)(\mu_1 + \varepsilon_{1,t}) + \theta_2(m, d)(\mu_2 + \varepsilon_{2,t}) + \theta_3(m, d)(\mu_3 + u_{m,t}), \quad (13.1)$$

where $\Delta \ln F_{m,t} = \ln F_{m,t} - \ln F_{m,t-1}$ is the change from $t-1$ to t of the log price of the futures contract that matures at m . $\varepsilon_t = (\varepsilon_{1,t} \varepsilon_{2,t})'$ is a vector of the latent factors that affect all contracts traded on t , with $E[\varepsilon_t] = 0$ and $E[\varepsilon_t \varepsilon_t'] = I_2 \forall t$, and $E[\varepsilon_t \varepsilon_s'] = 0$ for $t \neq s$. $u_{m,t}$ is the idiosyncratic error, representing shocks that are specific to the contract maturing at m . It is assumed that $E[u_{m,t}] = 0$ and $V[u_{m,t}] = 1 \forall m$ and t , and $E[u_{m,s} u_{n,t}] = 0$ for $s \neq t$ and/or $m \neq n$. That is, $u_{m,t}$ is uncorrelated serially and across concurrently traded contracts. $\theta_1(m, d)$ and $\theta_2(m, d)$ are the factor loadings determining the extent to which common shocks, $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$, are reflected in the price change of the contract maturing at m , and $\theta_3(m, d)$ determines the standard deviation of the shock specific to the contract maturing at m . In model (13.1), three coefficients μ_i are included to allow for a potentially non-zero deterministic change in the log futures price. They are multiplied by the associated $\theta_i(m, d)$ function so that they are interpreted as representing the forward premium associated with the two common factors ($i = 1, 2$) and idiosyncratic error ($i = 3$).†

The three terms, $\theta_i(m, d)$ for $i = 1, 2$ and 3 , are specified as deterministic functions of the contract delivery month (m) and time to delivery of the contract ($d = m - t$),

$$\theta_i(m, d) = \exp \left[a_{i,m,0} + a_{i,m,1}d + \sum_{k=1}^K \left(a_{i,m,2k} \sin \left(\frac{2\pi kd}{d_{\max}} \right) + a_{i,m,2k+1} \cos \left(\frac{2\pi kd}{d_{\max}} \right) \right) \right], \quad (13.2)$$

where $d_{\max}$ is the maximum days to maturity for which the model is estimated. Specification (13.2) allows the three terms $\theta_i(m, d)$ to be a flexible function of time-to-maturity (d) and permits this function to vary by contract delivery month (m).‡ The combination of m and d uniquely identifies a trade date (t) in the year through $d = m - t$. Thus, specification (13.2) also captures seasonal variation in the factor loadings and the idiosyncratic variance.

The unconditional variance of daily log futures returns is given by $V[\Delta \ln F_{m,t}] = \sum_{i=1}^3 \theta_i(m, d)^2$. Thus, the model can replicate very complex dynamics in the variance of log futures returns and its composition among the three components. Furthermore, the two latent factors affect all contract prices, whereas the idiosyncratic errors are uncorrelated across contracts. Therefore, correlations across concurrently

* The model as originally introduced by Smith (2005) is called the Partially Overlapping Time-Series or 'POTS' model as it was developed for the analysis of commodity futures return data, which usually form a partially overlapping time series or an unbalanced panel. In this paper, I refer to the model defined in (13.1)–(13.3) as a 'flexible model of futures returns' or simply 'flexible model' because the other two models examined in this paper are also applied to partially overlapping time-series data.

† The estimates of the three coefficients μ_i are very small for all four commodities examined in this paper. Restricting these coefficients to zero does not alter the results presented in Sections 13.3 and 13.4.

‡ The function $\theta_i(m, d)$ becomes more flexible with the number of trigonometric terms (K). Although this extra flexibility allows a better fit to the observed data, it also makes the model more sensitive to extreme observations. In the empirical estimation of the model in Section 13.3, I set $K = 3$ so that the model is flexible enough to capture seasonality and maturity effects while avoiding excess sensitivity to extreme observations.

traded contracts are determined by the share of the variance attributable to the two common factors. Specification (13.2) allows these cross-contract correlations to vary by season and across contracts.

For identification, the constraint is imposed as $a_{2,m,0} = -\sum_{k=1}^K a_{2,m,2k+1} - a_{2,m,1}d_{\max} - 10$ so that $\theta_2(m, d_{\max}) = 0$ for all m . That is, the loading of the second factor is close to zero at the maximum days to maturity. The condition is equivalent to that used by Schwartz and Smith (2000), which allows the two factors to be interpreted as representing the long-term (LT) and short-term (ST) factor.*

The conditional variance of latent factors ϵ_t is specified by a bivariate GARCH(1,1) model in a diagonal BEKK specification (Engle and Kroner 1985),

$$\begin{aligned} E[\epsilon_t \epsilon_t' | \mathfrak{I}^{t-1}] &= \mathbf{H}_t, \\ \mathbf{H}_t &= \mathbf{\Omega} + \beta \mathbf{H}_{t-1} \beta' + \alpha E[\epsilon_{t-1} \epsilon_{t-1}' | \mathfrak{I}^{t-1}] \alpha', \end{aligned} \quad (13.3)$$

where $E[\cdot | \mathfrak{I}^{t-1}]$ denotes that the expectation is conditional on the set of the information available at $t-1$, α and β are 2 by 2 diagonal matrices of parameters, and $\mathbf{\Omega}$ is a 2 by 2 diagonal matrix whose values are determined by $\mathbf{\Omega} = \mathbf{I}_2 - \beta \beta' - \alpha \alpha'$ due to $E[\epsilon_t \epsilon_t'] = \mathbf{I}_2$. This condition is required for the identification of the factor loadings $\theta_i(m, d)$ for $i = 1$ and 2.

Since the values of the two common factors are not observable, the coefficients in the model defined in (13.1)–(13.3) are estimated through the Kalman filter (Hamilton 1994). The measurement equation is formulated by stacking all prices observed on day t ,

$$\Delta \ln \mathbf{F}_t = \mathbf{C}_t + \Theta_{1,t} \epsilon_{1,t} + \Theta_{2,t} \epsilon_{2,t} + \Theta_{3,t} \mathbf{u}_t, \quad (13.1')$$

where $\Delta \ln \mathbf{F}_t$ is an n_t by one column vector of the daily log futures returns observed on day t with n_t representing the number of futures contracts traded on t , $\mathbf{C}_t$ is an n_t by one column vector of constants with its j th element given by $\sum_{i=1}^3 \theta_i(m_j, d_j) \mu_i$, $\Theta_{i,t}$ is an n_t -dimensional diagonal matrix with $\theta_i(m_j, d_j)$ on its j th diagonal element, and $\mathbf{u}_t$ is an n_t by one column vector with $u_{j,t}$ on its j th element.

The two state variables in ϵ_t are serially uncorrelated, whereas their conditional variances are serially related through the GARCH process in (13.3). The conditional expectation in the second equation of (13.3) is

$$E[\epsilon_{t-1} \epsilon_{t-1}' | \mathfrak{I}^{t-1}] = \epsilon_{t-1|t-1} \epsilon_{t-1|t-1}' + P_{t-1|t-1},$$

where $\epsilon_{t-1|t-1} = E[\epsilon_{t-1} | \mathfrak{I}^{t-1}]$ and $P_{t-1|t-1} = E[(\epsilon_{t-1} - \epsilon_{t-1|t-1})(\epsilon_{t-1} - \epsilon_{t-1|t-1})' | \mathfrak{I}^{t-1}]$, which are obtained through the Kalman filter.

The model defined in (13.1)–(13.3) follows an approach similar to the one introduced by Heath *et al.* (1992) for modelling the term structure of interest rates; they both specify the dynamics of futures prices (forward rates) directly by a set of common latent factors. Theoretically, these factor models fit perfectly to the observed futures price changes if they include as many latent factors as the number of futures prices observed per day and impose no restrictions on the values of factor loadings. In practice, empirical models include only a small number of factors and impose certain parametric structures on the factor loadings. For example, Cortazar and Schwartz (1994), in their analysis of commodity contingent claims, include only three factors and plot the factor loadings as a function of time-to-maturity only. These restrictions create a discrepancy between the model's implied and observed price movement.

* In the two-factor model of Schwartz and Smith (2000), the loading of the short-term factor in the futures price equation is given by $\exp(-\kappa\tau)$ where κ and τ represent the mean-reversion coefficient and time-to-maturity, respectively. The value of this loading decreases exponentially and converges to zero with τ , given $\kappa > 0$. The two-factor model of Sorensen (2002) reviewed in Section 13.2.2 shares this property.

Previous studies often do not model explicitly this residual component. The model defined in (13.1)–(13.3) differs from Cortazar and Schwartz (1994) in these regards; it specifies the two factor loadings as flexible parametric functions of time-to-maturity and contract delivery month. It also models parametrically the idiosyncratic error with its variance also specified by a flexible function. These specifications together allow the model to depict complex seasonal and cross-sectional variations in the variance and covariances of concurrently traded contract prices.

13.2.2 A conventional Two-Factor Term-Structure Model of Commodity Prices

A conventional approach to modelling a term structure of commodity prices is to specify the stochastic process of the spot price of the commodity and to derive from the stipulated spot price dynamics pricing formulas of futures and other derivative contracts. For example, the following two-factor Gaussian model is commonly considered for the analysis of various commodities with seasonality in demand and/or supply:^{*}

$$\begin{aligned}\ln S_t &= f(t) + x_t + z_t, \\ dx_t &= \mu dt + \sigma_x dw_x, \\ dz_t &= -\kappa z_t dt + \sigma_z dw_z, \\ dw_x dw_z &= \rho dt,\end{aligned}\tag{13.4}$$

where S_t is the spot price at period t , $f(t)$ is the seasonal mean price that is a deterministic function of time, x_t and z_t are the state variables representing, respectively, the long-term (LT) and short-term (ST) deviation from the seasonal mean price, dw_x and dw_z are the increments to the standard Brownian motion that are correlated through $dw_x dw_z = \rho dt$, and μ , κ , σ_x and σ_z are parameters determining, respectively, the drift rate, the mean-reversion rate, and the diffusion rate of the two stochastic factors.

The price in period t of the futures contract that matures at T is obtained as the period t conditional expectation, under the risk-neutral probability measure, of the spot price at T . It can be shown that, for the spot price following the process (13.4), the price of this futures contract is obtained as[†]

$$\ln F(t, T) = f(T) + A(\tau) + x_t + z_t e^{-\kappa\tau},\tag{13.5}$$

where

$$A(\tau) = \left(\mu - \lambda_x + \frac{\sigma_x^2}{2} \right) \tau + \frac{\sigma_z^2 (1 - e^{-2\kappa\tau})}{4\kappa} + \frac{(\rho\sigma_x\sigma_z - \lambda_z)(1 - e^{-\kappa\tau})}{\kappa},$$

$\tau = T - t$, is the time-to-maturity, and the two coefficients λ_x and λ_z represent the market prices of risk associated with the corresponding stochastic factor.

^{*} Model (13.4) has been considered, for example, for the analysis of electricity (Lucia and Schwartz 2002), natural gas (Manoliu and Tompaidis 2002) and agricultural commodity futures, such as corn, wheat and soybean (Sorensen 2002). Gibson and Schwartz (1990) and Nielsen and Schwartz (2004) also consider a two-factor model in analysing oil and copper, yet they parameterize the dynamics of two factors differently from (13.4) so that the two factors are interpreted as representing the spot price and convenience yield factor. Schwartz and Smith (2000) show that these models are identical to model (13.4) aside from the absence of seasonal variation in mean price.

[†] Futures price formula (13.5) assumes constant market prices of risk. See, for example, Sorensen (2002) for details.

The set of parameters defining model (13.4), $\Omega = \{\kappa, \sigma_x, \sigma_z, \lambda_x, \lambda_z, \rho\}$, is usually estimated with futures price data. To fit Equation (13.5) into multiple prices with different maturity dates observed per day, an error term is added to the right-hand side of (13.5), which makes the values of the two factors x and z unidentifiable. Thus, the model is commonly estimated with a filtering method. In state-space form, the model is presented as

$$\ln F(t, T_i) = f(T_i) + A(\tau_i) + x_t + z_t e^{-\kappa \tau_i} + u_{T_i,t} \quad (\text{measurement equation}), \quad (13.6)$$

$$\begin{aligned} z_t &= e^{-\kappa} z_{t-1} + v_{1,t}, \\ x_t &= \mu + x_{t-1} + v_{2,t} \quad (\text{transition equation}), \end{aligned} \quad (13.7)$$

where T_i is the maturity date of the i th contract ($i = 1, \dots, n_t$) observed on day t , $\mathbf{v}_t = (v_{1,t}, v_{2,t})'$ is serially uncorrelated and identically distributed with $\mathbf{v}_t \sim N(0, \mathbf{H})$, and $\mathbf{H}$ is the symmetric matrix with σ_x^2 and σ_z^2 on the main diagonal and $\rho\sigma_x\sigma_z$ off-diagonal. In (13.6), $u_{T_i,t}$ is the measurement error representing errors in reporting prices or factors affecting futures prices that are not accounted for by the two common factors. It is commonly assumed that $E[u_{T_i,t}] = 0 \forall t$ and T , and $E[u_{T_1,s}u_{T_2,t}] = 0$ for $s \neq t$ and/or $T_1 \neq T_2$. Thus, econometrically, $u_{T_i,t}$ represents the idiosyncratic error that is uncorrelated serially and contemporaneously across contracts. It is also commonly assumed that $V[u_{T_i,t}] = \sigma_T^2 \forall t$. That is, the variance is allowed to vary by the contract maturity date (T) but not by trade date (t) or time-to-maturity ($T - t$).

13.2.3 Model Comparison

A major difference between conventional term-structure models of commodity prices and the flexible model of futures returns defined in (13.1)–(13.3) is that the former specifies the dynamics of price *level*, whereas the latter specifies the dynamics of price *return*. By modelling price returns rather than level, the flexible model does not specify seasonality and other deterministic variations in the underlying spot price that result from demand/supply seasonality and other characteristics of the underlying commodity.* Thus, the model is free from bias in specifying such deterministic price variation.

To compare the two models in further detail, take the first difference of the futures price formula in (13.6),

$$\begin{aligned} \Delta \ln F_{T,t} &= \ln F_{T,t} - \ln F_{T,t-1} \\ &= B(\tau) + v_{2,t} + e^{-\kappa \tau} v_{1,t} + \Delta u_{T,t}, \end{aligned} \quad (13.8)$$

where

$$B(\tau) = \lambda_x + \left(\frac{\lambda_z - \rho\sigma_x\sigma_z}{\kappa} \right) e^{-\kappa \tau} (1 - e^{-\kappa}) - \frac{\sigma_x^2}{2} - \frac{\sigma_z^2 e^{-2\kappa}}{4\kappa} (1 - e^{-2\kappa \tau}),$$

and $\Delta u_{T,t} \sim N(0, 2\sigma_T^2)$ because $\Delta u_{T,t} = u_{T,t} - u_{T,t-1}$ and $u_{T,t} \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_T^2) \forall t$.

Comparison of (13.8) and (13.1) reveals three major benefits of the flexible model over the conventional approach in modelling a term structure of commodity prices with the specified spot price dynamics. First, model (13.1) specifies the factor loadings by flexible functions for both the LT and ST factors. In contrast, the factor loadings are determined by a small number of parameters defining the stochastic dynamics of the underlying state variables in conventional term-structure models. Specifically, for the two-factor model (13.8), the loading of the ST factor decreases exponentially with time-to-maturity at an identical rate for all contracts, whereas that of the LT factor is constant at unity for all contracts

* These deterministic price variations correspond to the seasonal mean price and deterministic trend (denoted as $f(T)$ and μ) in the conventional term-structure model (13.4). First differencing eliminates these terms and leaves only the innovation errors (v_1 and v_2) on the right-hand side of (13.8). The stochastic dynamics of the two state variables still remain in (13.8), yet only implicitly by restricting the functional forms of the factor loadings.

throughout the trading horizon. Second, the flexible model (13.1) specifies the variance of the idiosyncratic error by a flexible function of time-to-maturity and allows this function to vary across contract delivery months. In contrast, conventional term-structure models impose a simplistic structure on the variance of the measurement error $u_{i,t}$ with the variance allowed to vary only by the delivery month of contract but not by time-to-maturity. Third, the innovations to the state variables, $v_{i,t}$ ($i = 1, 2$), are specified to follow a bivariate GARCH process in the flexible model (13.1) while they are assumed homoskedastic in the conventional term-structure model (13.8).

Strong restrictions imposed on the stochastic dynamics of the underlying factors and the variance of the measurement error potentially lead conventional term-structure models to draw an erroneous portrait of price volatility of a storable commodity with demand and/or supply seasonality. In particular, the model in (13.6) and (13.8) stipulates that the variance attributable to the two common factors increases exponentially as the contract approaches maturity, whereas the variance of the measurement error does not vary with time-to-maturity.* Consequently, the model implies that correlation across concurrently traded contracts decreases monotonically with time-to-maturity. By contrast, the flexible model defined in (13.1)–(13.3) allows for the magnitude of futures price change resulting from the common market shocks and that resulting from contract-specific shocks to differ both by time-to-maturity and by contract delivery date. This flexibility allows the model to replicate the highly non-linear dynamics of commodity prices expected by the theory of storage. The model also gives the same flexibility to the variance structure of the two common factors and that of the idiosyncratic error and thus avoids the magnitude and dynamics of the cross-contract correlation being determined by the model specification.

In the next section, I estimate the flexible model of futures returns with empirical data from four markets and compare its estimation results with the estimate of the two alternative models. The first model is the conventional two-factor model defined in (13.4). I estimate the subset of the model parameters that appear in the model's first difference form (13.8), which is directly comparable to the flexible model. Since first differencing eliminates the deterministic variation in mean price level and deterministic variations in the two common factors, the comparison signifies the adequacy of the parsimonious specifications imposed in the conventional two-factor model on the stochastic dynamics of the underlying factors (as determinants of the factor loadings) and the variance of the measurement error. The second alternative model is the composite model in which factor loadings are specified as in the conventional two-factor model (hence imposing a restrictive specification on the stochastic dynamics of the latent factors), whereas the variance of the measurement error is specified by a flexible function as in (13.2),

$$\Delta \ln F_{T,t} = B(\tau) + v_{2,t} + e^{-\kappa\tau} v_{1,t} + \theta_3(m, d) u_{T,t}, \quad (13.9)$$

where $\theta_3(m, d)$ is as defined in (13.2). This model nests the two-factor model in (13.8). It is expected that the flexible specification of the variance of measurement error captures complex volatility dynamics, albeit partially.

13.3 Data and Estimation

This section empirically estimates the three models reviewed in [Section 13.2](#) with the unbalanced panel data from four commodity markets. The section starts with a description of the data examined and then reports results from estimating the three models with emphasis on their implied volatility dynamics.

* In other words, of the three components comprising the daily futures returns in model (13.8), the ST factor is the only component that can vary by time-to-maturity, and the measurement error is the only component that allows variation across contracts with different maturity dates.

13.3.1 Data

The three models presented in Section 13.2 are estimated with the data from the markets for the following four commodities with varying characteristics:

- natural gas—consumption good with strong seasonality in demand
- corn—consumption good with strong seasonality in supply
- crude oil—consumption good with very weak seasonality in demand and supply
- gold—investment good with virtually no seasonality either in demand or supply

The models are estimated using data on daily settlement prices of futures contracts traded at the New York Mercantile Exchange (crude oil, natural gas and gold) and Chicago Board of Trade (corn). The data analysed in this paper are from the period between 1984/1/1 and 2007/12/31 for corn and gold, 1984/4/1 and 2007/12/31 for crude oil, and 1991/4/1 and 2007/12/31 for natural gas. For each contract, daily prices to the last trading day of the contract are used for analysis.* Since long-dated contracts do not trade actively, contracts of more than 12 months to maturity are excluded from the analysis, except that contracts up to 18 months to maturity are analysed for corn.† Excluding these observations leaves 70,800 prices among 307 contracts for crude oil, 52,780 prices among 223 contracts for natural gas, 43,820 prices among 168 contracts for gold, and 48,762 prices among 142 contracts for corn.

All three models are estimated by the method of maximum likelihood with the likelihood obtained through the Kalman filter as described in Section 13.2. I first estimate the conventional two-factor model (13.8), which is the most parsimonious of the three models. The composite model (13.9) differs from model (13.8) only in the specification of the variance of the measurement error. The model is estimated with the starting values of the coefficient vector $\mathbf{a}_3$ that minimize the sum of the squared differences between $\theta_3(m, d; \mathbf{a}_3)^2$ and the squared residuals from the estimated model (13.8). Finally, for the flexible model (13.1), I obtain the starting values in two steps: (i) calculate the variance attributable to each of the three components (LT and ST factor, and the idiosyncratic error) from the estimated composite model and (ii) find the values of each coefficient vector $\mathbf{a}_i (i = 1, \dots, 3)$ that minimize the sum of the squared differences between $\theta_i(m, d; \mathbf{a}_i)$ and the predicted values of the corresponding component in the composite model calculated in step (i). Robustness is checked by estimating the model with different sets of starting values obtained by distributing a fraction of the variance of the idiosyncratic error from the estimated composite model to the other two components in step (i) of the above two-step procedure.

13.3.2 Estimation Results: Model Specification

Table 13.1 summarizes the results from the specification test. It shows that, for all four commodities considered, the flexible model of futures returns is preferred to the other two models, and the composite model is preferred to the conventional two-factor model according to both the Akaike and Schwarz Information Criteria. The results provide strong evidence that the conventional two-factor model stipulating restrictive structures on the factor loadings and the variance of the measurement error is not supported empirically for all four commodities. Surprisingly, the model is not supported even for gold for which the storage cost is not significant and virtually no demand or supply seasonality exists.

* Crude oil and natural gas contracts cease trading before the delivery month, whereas corn and gold contracts trade into the delivery month. See the exchange's website (www.cmegroup.com) for details on contract specifications.

† Corn exhibits strong supply seasonality with the harvest usually arriving around September to early November. The theory of storage suggests that the inter-temporal price link breaks at the end of the crop year, creating potentially very complex price dynamics for the September and December contracts at around 12 months before maturity. I analyse the contracts as far as 18 months to maturity to capture potentially very interesting price movements of these two contracts in this period.


TABLE 13.1 Model Selection Test

	Commodity			
	Corn	Crude Oil	Gold	Natural Gas
Sample size				
<i>All</i>	48762	70800	43820	52780
<i>By contract delivery month</i>				
Jan		5970		4408
Feb		5882	7255	4411
Mar	9988	5923		4406
Apr		5904	7296	4383
May	9517	5864		4369
June		5853	7312	4391
July	9764	5811		4400
Aug		5803	7328	4402
Sep	9255	5887		4403
Oct		5919	7334	4402
Nov		5933		4404
Dec	10238	6051	7295	4401
<i>By days to maturity (business days before the first day of delivery month)</i>				
–21 – 0	1728		3057	
1 – 42	5665	10421	6800	8238
43 – 84	5665	12463	6797	8930
85 – 126	5664	12460	6794	8925
127 – 168	5665	12270	6795	8926
169 – 210	5664	12025	6791	8920
211 – 252	5662	11161	6786	8841
253 – 294	5570			
295 – 336	4756			
337 – 378	2723			
Number of contracts	142	307	168	223
Number of trading days	6808	6204	6779	4441
Number of parameters estimated				
Flexible	124	285	147	285
Conventional two factor	11	18	12	18
Composite	46	102	54	102
Akaike Information Criterion				
Flexible	–428928	–697588	–558094	–431278
Conventional two factor	–400078	–618955	–525549	–393527
Composite	–416075	–691672	–551946	–421183
Schwarz Information Criterion				
Flexible	–427837	–694975	–556817	–428749
Conventional two factor	–399981	–618790	–525444	–393367
Composite	–415670	–690737	–551477	–420277


13.3.3 Estimation Results: Flexible Model

Figures 13.1 through 13.4 illustrate the results of estimating the flexible model of futures returns for the four commodities. These figures plot, for each contract delivery month: (a) the estimated loadings of the LT factor, (b) the loadings of the ST factor, (c) the standard deviation of the idiosyncratic error and (d) the share of the total variance accounted for by the two common factors, which are all aligned by trade date. These components are calculated as $\theta_i(m, d; \hat{\mathbf{a}}_{i,m})$ for the first three components ($i = 1, 2$ and 3 , respectively, for components (a), (b) and (c))

(a) Factor loading 1



(b) Factor loading 2



(c) Standard deviation of idiosyncratic error






FIGURE 13.1 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Natural gas.

(Continued)


(d) Share of total variance accounted for by two common factors

**FIGURE 13.1 (Continued)** Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Natural gas.


(a) Factor loading 1



(b) Factor loading 2

**FIGURE 13.2** Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Corn.*(Continued)*

(c) Standard deviation of idiosyncratic error



(d) Share of total variance accounted for by two common factors




FIGURE 13.2 (Continued) Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Corn.

(a) Factor loading 1






FIGURE 13.3 Factor loadings, variance of idiosyncratic error, and share of total variance accounted for by two common factors in flexible model: Crude oil.

(Continued)

(b) Factor loading 2



(c) Standard deviation of idiosyncratic error



(d) Share of total variance accounted for by two common factors






FIGURE 13.3 (Continued) Factor loadings, variance of idiosyncratic error, and share of total variance accounted for by two common factors in flexible model: Crude oil.

(a) Factor loading 1



(b) Factor loading 2



(c) Standard deviation of idiosyncratic error




FIGURE 13.4 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Gold.

(Continued)

(d) Share of total variance accounted for by two common factors




FIGURE 13.4 (Continued) Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in flexible model: Gold.

and $\sum_{i=1}^2 \theta_i(m, d; \hat{\mathbf{a}}_{i,m})^2 / \sum_{i=1}^3 \theta_i(m, d; \hat{\mathbf{a}}_{i,m})^2$ for component (d), where $\hat{\mathbf{a}}_{i,m} = \{\hat{a}_{i,m,0}, \dots, \hat{a}_{i,m,2K}\}$ is the vector of coefficients estimated for each of the three components ($i = 1, 2, 3$), each delivery month (m), and for each of the four commodities.*

(a) Natural gas

In Figure 13.1(a), the loadings of the LT factor estimated for natural gas indicate two notable features. First, for all 12 contracts, the estimated factor loading increases as the contract approaches the maturity date. Second, the factor loadings in the last few months of trading are substantially higher for the contracts maturing in winter than those maturing in summer. The loadings of the ST factor exhibit the same features, but in greater magnitude than those observed for the LT factor (Figure 13.1(b)). In addition, for all 12 contracts, the loading of the ST factor starts increasing rapidly in May, before which it is virtually zero. In Figure 13.1(c), the variance of the idiosyncratic error, particularly that for winter contracts, increases very rapidly as the contract approaches maturity. This indicates that high volatility in the last one month of trading, commonly referred to as the maturity effect, represents market shocks that are specific to each contract and are of a very short-term nature. Unlike the two common factors, the idiosyncratic errors are not contemporaneously correlated across concurrently traded contracts. Thus, a rapid increase in the variance of the idiosyncratic error implies that correlations between nearby and distant futures contracts decrease rapidly over the winter season (Figure 13.1(d)).

These estimates of the factor loadings and the variance of the idiosyncratic error in the estimated flexible model are consistent with the price dynamics implied by the theory of storage for natural gas. In the flexible model (13.1), the total variance of the log futures price change is given by $\sum_{i=1}^3 \theta_i(m, d; \hat{\mathbf{a}}_{i,m})^2$ at d days before the contract maturing in the month m . Thus, high factor loadings and high variance of the idiosyncratic error translate into high volatility of winter contracts in the last few months of trading. During this peak-demand period, tight demand–supply imbalance causes demand or supply shock of even a small magnitude to follow a large price swing, which cannot be absorbed through adjusting inventory because the inventory is effectively zero at the end of the demand season. Low inventory

* The denominator in the formula for component (d) represents the total variance, owing to the assumption that the two latent factors and the idiosyncratic error are uncorrelated.

also means that the inter-temporal price linkage breaks at the end of the winter peak-demand season because, in any normal year, no physical stock is carried over from late winter (when price peaks) to early spring (when price is the lowest). The estimated flexible model reflects this feature with a large share of price variation accounted for by the ST factor and the idiosyncratic error for the December through March contracts in their last few months of trading. During the same period, the loadings of the ST factor and the variance of the idiosyncratic error stay very low for contracts maturing in May and thereafter.

(b) Corn

In [Figure 13.2](#), the estimated flexible model reveals complex volatility dynamics for corn. The depicted volatility pattern differs from natural gas, yet it is characterized by seasonality in the supply of the underlying commodity. In [Figures 13.2\(a\) and 13.2\(b\)](#), the loadings of both the LT and ST factor start increasing for all five contracts around April and peak in July to August. This observation implies that large price fluctuations during this period are highly correlated across the five contracts maturing in the post-harvest season. High price volatility during this period reflects the arrival of important information. In particular, corn crops in the U.S. are typically planted in early April through June and harvested later in the year, usually from September to early November. The date and yield of harvest are determined by weather conditions during summer. Thus, the contracts maturing post-harvest exhibit large price fluctuations and high cross-contract price correlations during this period. Volatility starts decreasing in mid-summer after most weather conditions are revealed and reaches the lowest point when actual harvesting is realized around October. Over these periods, prices move very closely for five contracts maturing post-harvest, because corn is an annual crop and one harvest in the current year needs to be stored for consumption until the new harvest arrives in the subsequent year. In normal years, the current harvest is fully consumed during the demand season and no inventory is carried over to the post-harvest season. Thus, the contracts maturing post-harvest show minimal price movement before these crops are planted in early spring. [Figure 13.2\(c\)](#) indicates that much of the high volatility in the last one month of trading originates in the contract-specific errors. The variance of the idiosyncratic error is particularly high for the July contract. This can be explained by low inventory at the end of the demand year, which does not allow unexpected demand and/or supply shocks to be absorbed through inventory adjustment. High variance of contract-specific shock means that a small share of variance is accounted for by the two common factors. This share and, consequently, the correlation among concurrently traded contracts decrease rapidly in the last one month of the trading period, particularly for the July contract ([Figure 13.2\(d\)](#)).

(c) Crude oil

Crude oil is generally thought to have much weaker demand seasonality than natural gas. However, the estimated flexible model indicates a moderate seasonality and maturity effect for the volatility of the crude oil price. In [Figure 13.3\(a\)](#), the estimated loading of the LT factor is slightly higher for all 12 contracts during winter months. The volatility also increases for all 12 contracts in the last two months of trading, and volatility in this period is slightly higher for winter (January and February) and summer (July and August) contracts. Much of this high volatility in the last two months of trading is captured by the ST factor and the idiosyncratic error ([Figures 13.3\(b\) and 13.3\(c\)](#)), indicating that this high volatility is caused by shocks that are not persistent and have little impact on the prices of distant maturity contracts. The variance of the idiosyncratic error increases rapidly in the last month of trading and is particularly high for the January and two summer contracts, causing substantial declines in correlations between these contracts and other concurrently traded contracts ([Figure 13.3\(d\)](#)).

(d) Gold

In Figure 13.4, the estimated flexible model shows virtually no seasonality or maturity effect for gold price volatility. For all six contracts, the estimated loadings of the LT factor are slightly higher from October to March, yet this seasonal difference is negligible, with variation ranging no more than 3.5% of the average (Figure 13.4(a)). Another feature that differentiates gold from the other three commodities is that the estimated loadings of the ST factor and the variance of the idiosyncratic error are very small for all contracts; they are on average 5.7% and 1.7%, respectively, of the estimated loadings of the LT factor. These results imply that much of the price shocks to the gold markets are of a long-term nature. High persistence of price shocks is as expected because gold is traded primarily for investment rather than consumption and, unlike other exhaustible resources, the amount of its deposits is well known. Due to large loadings of the LT factor, relative to the variance of the idiosyncratic error, the two common factors account for a large share of the total variance, resulting in high cross-contract correlations for gold as seen in Figure 13.4(d).


13.3.4 Estimation Results: Conventional Two-Factor Model

Figures 13.5 through 13.8 present the results of estimating the conventional two-factor model in the first difference form (13.8) for the four commodities. These figures plot, for each contract: (a) the variance of futures price attributable to the two common factors, which is calculated as, $\hat{\sigma}_x^2 + \hat{\sigma}_z^2 e^{-2\hat{\kappa}\tau} + \hat{\rho}\hat{\sigma}_x\hat{\sigma}_z e^{-\hat{\kappa}\tau}$, (b) the variance of the idiosyncratic error $\hat{\sigma}_m^2$ and (c) the share of the variance accounted for by the two common factors as calculated by $1 - \hat{\sigma}_m^2 (\hat{\sigma}_x^2 + \hat{\sigma}_z^2 e^{-2\hat{\kappa}\tau} + \hat{\rho}\hat{\sigma}_x\hat{\sigma}_z e^{-\hat{\kappa}\tau} + \hat{\sigma}_m^2)^{-1}$. In panel (a) of Figures 13.5 through 13.8, the variance attributable to the two common factors increases exponentially as the contract approaches maturity for all four commodities. This property, as discussed in Section 13.2, is the direct result of the model specification that the LT and ST factors follow the GBM and MR processes, respectively.


The conventional model allows the variance of the idiosyncratic error to vary by contract delivery date but not by time-to-maturity. In Figure 13.5(b), the model estimated for natural gas indicates a higher variance of the idiosyncratic error for winter (January through March) contracts than for the other contracts. However, the estimated variance of the idiosyncratic error, even for these winter contracts, is negligible in size when compared with the variance attributable to the two common factors. These estimates cause the model to imply very high cross-contract correlation, which increases as contracts approach maturity because the variance attributable to the ST factor increases exponentially. These implications for the magnitude of the cross-contract correlation and its dynamics over time-to-maturity are exactly opposite to the implications of the flexible model.

These results signify the severity of the restrictions imposed by the specifications of the conventional two-factor model. Of the three stochastic components on the right-hand side of (13.8), the ST factor is the only component that captures the dynamics of the price variance over the trading horizon, yet it only allows the variance to decrease exponentially with time-to-maturity at an identical rate for all contracts. It neither permits non-monotonic change of the price variance, nor allows the variance to change at different rates across contracts. The model also restricts the idiosyncratic error to be the only component to capture the cross-contract difference in price variance. The estimation results described above are determined by these specifications of the model. For natural gas, the model allocates a large share of the price variance to the ST factor to capture the strong maturity effect. It implies low variance of the idiosyncratic error because seasonal variation in the volatility of the natural gas price is small relative to the maturity effect. The severity of these restrictions is apparent when the model is compared with the flexible model, in which the three components are equally flexible in their seasonal and cross-sectional variation and the composition of the observed price movements among the three components is determined by cross-contract price correlations.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors




FIGURE 13.5 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in conventional two-factor model: Natural gas.

In Figures 13.6 through 13.8, the conventional two-factor model estimated for corn, crude oil and gold exhibits the same results as for natural gas. For all three commodities, the variance attributable to the two common factors decreases gradually with time-to-maturity at the identical rate for all contracts. The variance of the idiosyncratic error exhibits cross-contract variation, with variance slightly higher for the July corn contract, the August gold contract and the winter crude oil contracts. However, for all three commodities and for all contracts, the variance of the idiosyncratic error is constant over the trading horizon and much smaller than the variance attributable to the two common factors. The results imply that the two common factors account for a large share of price variance and that this share increases as the contract approaches maturity; the implications are exactly opposite to the estimated flexible model. For gold only, the difference between the flexible model and the conventional two-factor model is small because the very weak seasonality and maturity effect leads both models to assign a dominant share of the price variance to the LT factor, resulting in very small loading of the ST factor and a small variance of the idiosyncratic error.

13.3.5 Estimation Results: Composite Model


Figures 13.9 through 13.12 show the results from estimating the composite model (13.9).^{*} In Figure 13.9(a), the variance attributable to the two common factors for natural gas exhibits the same dynamic pattern but is slightly smaller in magnitude than the estimate in the conventional two-factor model. In Figure 13.9(b), the variance of the idiosyncratic error estimated for the composite model is substantially greater than the estimates for the conventional two-factor model and exhibits a strong seasonality and maturity effect. The strong seasonality and maturity effect of the idiosyncratic error is reflected in the total variance, resulting in a dynamics similar to that estimated for the flexible model. Nonetheless, the share of the total variance accounted for by the two common factors exhibits different dynamics between the two models simply because the compositions of the price variance among the three components differ between the two models (see Figure 13.9(c) and Figure 13.1(d)).

Figures 13.10 through 13.12 show that the composite model estimated for corn, crude oil and gold yields results similar to those for natural gas. The variance attributable to the two common factors decreases with time-to-maturity yet at a slower rate than in the conventional two-factor model. On the other hand, the variance of the idiosyncratic error is estimated substantially greater for the composite model than for the conventional two-factor model. The variance estimate for the composite model also exhibits a strong maturity effect with variance in the last one month of trading being particularly large for the July contract for corn and the January, July and August contracts for crude oil. The maturity effect captured by the idiosyncratic error dominates the increase in the loading of the ST factor, implying that cross-contract correlation decreases rapidly in the last few months of trading. The increase in the variance of the idiosyncratic error near the maturity date is much smaller for gold, resulting in only a marginal reduction in the cross-contract correlation near the maturity dates.


In summary, estimates of the flexible model of futures returns reveal that the volatility of the three commodity futures, natural gas, corn and crude oil, exhibits a strong seasonality and maturity effect. For natural gas, the volatility in the last few months of trading is particularly high in winter months when demand peaks and inventory is low. Much of the high volatility during this period is captured by the ST factor and the idiosyncratic error because low inventory breaks the inter-temporal price link, resulting in low cross-contract price correlation. For corn, volatility increases in spring through early summer when crops are planted and as much of the weather shocks affecting the growth of these crops are revealed. During this period, contracts maturing in the post-harvest season exhibit high price

^{*} In these figures, the variance of futures price attributable to the two common factors is calculated in the same way as for the conventional two-factor model, whereas the variance attributable to the idiosyncratic error is given as $\theta_3^2(m, d; \hat{\mathbf{a}}_{3,m})$.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors






FIGURE 13.6 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in conventional two-factor model: Corn.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors






FIGURE 13.7 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in conventional two-factor model: Crude oil.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for two common factors






FIGURE 13.8 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in conventional two-factor model: Gold.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors






FIGURE 13.9 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in composite model: Natural gas.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors






FIGURE 13.10 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in composite model: Corn.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors






FIGURE 13.11 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in composite model: Crude oil.

(a) Variance attributable to two common factors



(b) Variance attributable to idiosyncratic error



(c) Share of variance accounted for by two common factors




FIGURE 13.12 Factor loadings, variance of idiosyncratic error and share of total variance accounted for by two common factors in composite model: Gold.

correlation. In the same period, the July contract that matures before harvest is subject to high contract-specific shock because low inventory at the end of the demand season breaks the inter-temporal price link. Crude oil exhibits high price volatility both in winter and summer months, but the seasonality is moderate relative to natural gas and corn.

The conventional two-factor model fails to capture these complex volatility dynamics of the three consumption commodities because of strong restrictions imposed on the factor loadings and the variance of the idiosyncratic error. In particular, the model allows neither a cross-contract nor a seasonal variation in the factor loadings. Neither does it allow a variation in the variance of the idiosyncratic error by time-to-maturity. Consequently, the model captures cross-contract variation in the price volatility solely by the idiosyncratic error and the maturity effect by the ST factor only. Flexible specification of the variance of the idiosyncratic error permits the composite model to alleviate, albeit imperfectly, the restrictions imposed in the conventional two-factor model. The model replicates reasonably well the complex volatility dynamics of the consumption commodities, yet it implies dynamics of cross-contract correlations that are substantially different from the pattern implied by the flexible model.

13.4 Implications for the Optimal Hedging Strategy

In the previous section, estimates of the factor loadings and the variance of the idiosyncratic error differ substantially across the three models. This section compares the three models with respect to their implications for an optimal hedging strategy.

13.4.1 Optimal Hedging

I consider a trader with a spot position Q in period t , who simultaneously takes a short position in X futures contracts for delivery at $\tau > t$. At $t + k < \tau$, the trader clears his position by selling Q units in the spot market and buying X futures contracts for delivery at τ .

The return to this trader's portfolio from t to $t + k$ is

$$\begin{aligned} W_{t+k} &= \Delta \ln S_{t+k} Q - \Delta \ln F_{\tau, t+k} X \\ &= (\Delta \ln S_{t+k} - \eta \Delta \ln F_{\tau, t+k}) Q, \end{aligned}$$

where $\Delta \ln S_{t+k}$ and $\Delta \ln F_{\tau, t+k}$ are, respectively, the change from t to $t + k$ of the log spot price and the log price of the futures contract for delivery at τ , and $\eta = XQ^{-1}$ is the hedge ratio. The variance of the portfolio return is

$$V[W_{t+k}] = Q^2 (V[\Delta \ln S_{t+k}] + \eta^2 V[\Delta \ln F_{\tau, t+k}] - 2\eta \text{cov}[\Delta \ln S_{t+k}, \Delta \ln F_{\tau, t+k}]),$$

which is minimized when the hedge ratio is set as

$$\eta^* = \frac{\text{cov}[\Delta \ln S_{t+k}, \Delta \ln F_{\tau, t+k}]}{V[\Delta \ln F_{\tau, t+k}]} \quad (13.10)$$

The minimum variance attained by this hedge ratio is

$$V[W_{t+k} | \eta^*] = V[\Delta \ln S_{t+k}] (1 - \rho_{\tau, t+k}^2) Q^2, \quad (13.11)$$

where $\rho_{\tau,t+k}$ is the correlation between the change in the log spot price and the change in the log price of the futures contract for delivery at τ over the period between t and $t + k$.

In the above variance minimization problem, the delivery date τ of the futures contract included in the portfolio is taken as exogenous. In many organized exchanges, however, multiple contracts with different maturity dates are traded simultaneously. Thus, the hedger can choose from these multiple contracts one that attains the minimum portfolio variance. Equation (13.11) indicates that, given the time of entry (t) and hedging horizon (k), the portfolio variance is minimized when it includes the contract that exhibits the highest correlation with the spot price over the hedging horizon. Once this contract is identified, the optimal hedge ratio is determined according to (13.10) as the ratio of the covariance between the spot and futures returns to the variance of futures returns.

13.4.2 Optimal Hedging Strategy According to the Three Models

It is apparent from (13.10) and (13.11) that the specifications of the variance and covariance dynamics of spot and futures prices play an important role in determining the optimal hedging strategy. Given the time of entry (t) and hedging horizon (k), one can calculate the optimal hedge ratio (η) and the associated portfolio variance for every possible choice of futures contracts included into the portfolio (τ), based on the three models of daily futures returns estimated in [section 13.3](#). In particular, the variance of log futures returns in the denominator of expression (13.10) is calculated as, for the conventional two-factor model,

$$V[\Delta \ln F_{\tau,t}] = \hat{\sigma}_x^2 + \hat{\sigma}_z^2 e^{-2\hat{\kappa}d} + 2\hat{\rho}\hat{\sigma}_x\hat{\sigma}_z e^{-\hat{\kappa}d} + \hat{\sigma}_\tau^2, \quad (13.12a)$$

for the flexible model,

$$V[\Delta \ln F_{\tau,t}] = \sum_{i=1}^3 \theta_i^2(\tau, d; \hat{\mathbf{a}}_i), \quad (13.12b)$$

and for the composite model,

$$V[\Delta \ln F_{\tau,t}] = \hat{\sigma}_x^2 + \hat{\sigma}_z^2 e^{-2\hat{\kappa}d} + 2\hat{\rho}\hat{\sigma}_x\hat{\sigma}_z e^{-\hat{\kappa}d} + \theta_3^2(\tau, d; \hat{\mathbf{a}}_3) \quad (13.12c)$$

where $d = \tau - t$ is the time to delivery of the contract maturing at τ . Similarly, using the nearby futures price as a proxy for the spot price,* one can calculate the covariance between the log returns to the spot (nearby futures) and futures contract for delivery at τ as, for the conventional two-factor and composite model,

$$\text{cov}[\Delta \ln S_t, \Delta \ln F_{\tau,t}] = \hat{\sigma}_x^2 + \hat{\sigma}_z^2 e^{-\hat{\kappa}(d_0+d)} + \hat{\rho}\hat{\sigma}_x\hat{\sigma}_z (e^{-\hat{\kappa}d_0} + e^{-\hat{\kappa}d}) \quad (13.13a)$$

and, for the flexible model,

$$\text{cov}[\Delta \ln S_t, \Delta \ln F_{\tau,t}] = \sum_{i=1}^2 \theta_i(\tau_0, d_0; \hat{\mathbf{a}}_i) \theta_i(\tau, d; \hat{\mathbf{a}}_i) \quad (13.13b)$$

where $d_0 = \tau_0 - t$ and τ_0 is the delivery date of the nearby contract.

* This approximation is justified by the standard arbitrage argument that the price of the futures contract converges to the spot price as the contract approaches the maturity date (see, for example, Hull (2002)).

The first expression in (13.13) indicates that, for the conventional two-factor model and the composite model, the covariance between the log spot and log futures returns declines monotonically with time-to-maturity (d) of the futures contract, which results from the specification that the two common factors follow the GBM and MR processes in (13.4). Furthermore, for the conventional two-factor model, the variance of log futures returns also declines monotonically with time-to-maturity in (13.12) yet at a slower pace than the decline in the covariance of the spot and futures returns. Thus, the correlation between the two prices declines monotonically with time-to-maturity and the model proposes that the hedger should include in the portfolio the contract that is the closest to maturity to minimize the portfolio variance. This implication for the optimal hedging strategy generally applies to the composite model except that the model suggests the use of a more distant contract when the close-to-maturity contracts are increasingly subject to high idiosyncratic variance $\theta_3(\tau, d; \hat{\mathbf{a}}_3)$ (in which case, the variance of the futures contract increases faster than the covariance between the two prices). For the flexible model, the estimated factor loadings and the variance of the idiosyncratic error exhibit substantial seasonal and cross-contract variations. Because the resulting spot-futures correlations also exhibit seasonal and cross-sectional variations, the model proposes potentially very complex hedging strategies that require both the size (η) and position (T) in the futures market to be adjusted frequently across seasons.


Using the variance and covariance implied by each of the three models, I calculate the optimal hedge ratio for a one-day hedging horizon ($k = 1$) and t ranging from the first day to the last day of a calendar year.* For each t , I compare the futures contracts maturing in the subsequent 12 months by their correlation with the spot price and calculate the optimal hedge ratio for the portfolio including the contract that exhibits the highest correlation with the spot price.

Figure 13.13 illustrates, for each of the four commodities and for each of the three models, how the futures contract included in the optimal portfolio shifts by the date of entry. In the figure, the optimal contract differs substantially among the three models. The optimal hedging strategy based on the conventional two-factor model includes the second position contract which is the closest to maturity after the nearby contract. This result is expected because the model stipulates that, for each contract, the share of the variance accounted for by the two common factors increases as the contract approaches its maturity date. Thus, on any particular day, the contract closer to maturity exhibits a higher correlation with the spot price unless it receives an exceptionally large idiosyncratic variance. The exception is observed for corn in July through mid-August, during which the nearby (September) contract exhibits higher correlation with the third position (March) contract than with the second position (December) contract because of the high idiosyncratic variance of the latter.


The flexible model and composite model also suggest the use of the near-to-maturity (second position) contract to hedge against the spot price risk for gold and corn. For crude oil and natural gas, the two models suggest the use of a more distant contract, typically the third position contract for crude oil and the third or fourth position contract for natural gas. The two models suggest the use of distant contracts because the contracts closer to maturity are subject to high idiosyncratic shocks and hence exhibit low correlation with the nearby contract. For natural gas, the optimal portfolio includes the June or July contract early in the calendar year (February to May). During this time, contracts for earlier maturity are subject to high idiosyncratic shocks because low inventory at the end of the demand year disconnects the inter-temporal price links. The optimal contract shifts gradually from June to mid-August and, after that, it switches to a winter (December or January) contract. In this post-summer season, winter contracts are subject to little idiosyncratic shocks and much of their price movements

* I consider a hedging horizon of a single day because, in the absence of a transaction cost for adjusting the futures position, the best (variance-minimizing) strategy for a longer hedging horizon is to alter the choice of contract (τ) and the hedge ratio (η) every day so as to follow the optimal strategy for a one-day horizon. Staying with the previous-day position could be optimal in the presence of transaction cost. The best strategy in such a setting will depend on various factors such as the length of the hedging horizon, the size of the transaction cost and the risk-aversion coefficient of the hedger. This is beyond the scope of the paper and hence is left for future research.


(a) Crude oil



(b) Natural gas



(c) Gold



(d) Corn




FIGURE 13.13 Position of contract included into the variance minimizing portfolio relative to the spot (nearby futures) position.

reflect common market shocks. Around the beginning of November, winter contracts start receiving idiosyncratic shocks and the optimal contract is replaced with a more distant contract.

Figure 13.14 plots the optimal hedge ratio against the time of entry for each of the four commodities. In Figure 13.14(c), the optimal hedge ratio for gold is close to one throughout the year for all three models. For gold, the variance of the idiosyncratic error is very small and a large share of the price variance is accounted for by the LT factor in all three models. Thus, cross-contract correlation is very high and the variance and covariance in expression (13.10) are about the same magnitude, resulting in the optimal hedge ratio of close to one. Similarly, the optimal hedge ratio for corn remains close to one in Figure 13.14(d), yet it is slightly lower for the flexible model than for the other two models because the model implies higher idiosyncratic variances (and hence a lower covariance of the two prices) than the other two models.

For the other two commodities, the optimal hedge ratio differs substantially across the three models. For both natural gas and crude oil, the optimal hedge ratio suggested by the conventional two-factor model shows little seasonal variation while it increases gradually towards the end of each calendar month. The model stipulates that the variance attributable to the two common factors increases exponentially as the contract approaches its maturity date. The specification follows that the covariance between spot and futures returns increases faster than the variance of futures return, resulting in a gradual increase in the optimal hedge ratio.

The composite model implies a similar strategy as the conventional two-factor model for hedging crude oil price risk, except that it suggests a much smaller variation of the optimal hedge ratio within each month. In contrast, the strategy based on the flexible model indicates a large increase in the optimal hedge ratio within each month simply because the estimated factor loadings increase much faster in this model than in the other two models. For natural gas, the optimal hedge ratio suggested by the flexible and composite model exhibits substantial seasonal variation, reflecting the seasonal variation in the position of the futures contract included in the optimal portfolio. In general, the optimal hedge ratio is high when the portfolio includes a distant contract, because a distant contract is subject to little idiosyncratic variance and hence results in a low variance of futures return in the denominator of (13.10). An optimal hedge ratio substantially above unity is often suggested because the futures contract included in the portfolio exhibits small variation, a dominant share of which represents common market shocks.

Table 13.2 presents the sum of the portfolio variance over a one-year hedging horizon expressed as the ratio to the variance of the log spot price without hedging.* In the table, all three hedging strategies reduce portfolio variance dramatically. The three strategies are particularly effective in reducing spot price risk for gold, for which cross-contract correlation is particularly high throughout the year.

* These numbers are calculated under the assumption that the futures returns follow the process as specified in the estimated flexible model, which receives the strongest empirical support for all four commodities. Specifically, the variance of the optimal portfolio is calculated, for each of the three hedging strategies and for each entry date (t), as

$$V[W_{t+1}] = V[\Delta \ln S_{t+1}] + \eta^*(t)^2 V[\Delta \ln F_{\tau^*(t), t+1}] - 2\eta^*(t) \text{cov}[\Delta \ln S_{t+1}, \Delta \ln F_{\tau^*(t), t+1}],$$

where $\tau^*(t)$ and $\eta^*(t)$ are the delivery month of the futures contract included in the portfolio and the hedge ratio according to the optimal strategy as seen in Figures 13.13 and 13.14, respectively, and the variance and covariance are evaluated according to the estimated flexible model, i.e.


$$V[\Delta \ln S_{t+1}] = \sum_{i=1}^3 \theta_i^2(\tau_0, d_0; \hat{\mathbf{a}}_i),$$

$$V[\Delta \ln F_{\tau^*(t), t+1}] = \sum_{i=1}^3 \theta_i^2(\tau^*(t), d^*(t); \hat{\mathbf{a}}_i), \text{ and}$$


$$\text{cov}[\Delta \ln S_{t+1}, \Delta \ln F_{\tau^*(t), t+1}] = \sum_{i=1}^2 \theta_i(\tau_0, d_0; \hat{\mathbf{a}}_i) \theta_i(\tau^*(t), d^*(t); \hat{\mathbf{a}}_i),$$

where $d^*(t) = \tau^*(t) - t$.


(a) Crude oil



(b) Natural gas



(c) Gold



(d) Corn


**FIGURE 13.14** Optimal hedge ratio.

TABLE 13.2 Theoretical Valuation of Hedging Performance

	Conventional	Flexible	Composite
Crude oil	20.94%	18.54%	18.91%
Natural gas	26.18%	20.09%	25.54%
Gold	0.14%	0.14%	0.17%
Corn	8.78%	7.97%	9.49%

Portfolio variance is relatively high for crude oil and natural gas, ranging from around 18 to 26% of the variance of the spot price. Hedging strategies are less effective for these two commodities because the nearby contract is subject to a large contract-specific shock near maturity.

For all four commodities, portfolio variance is the smallest for a strategy based on the flexible model. The composite model performs better than the conventional model in hedging spot price risks for crude oil and natural gas. However, the strategy results in a substantially greater portfolio variance than the strategy based on the flexible model for natural gas. This result signifies the importance of specifying factor loadings by a flexible functional form to properly account for not only seasonal and temporal variation in the price variance, but also the variation in the cross-contract correlation, the latter of which plays a significant role in designing an effective hedging strategy.

13.5 Conclusion

In this study, I compare a conventional two-factor term-structure model of commodity futures with an alternative modelling approach of specifying the variance of daily futures returns directly by flexible functions. Empirical estimation of the flexible model of futures returns with daily futures price data from four commodity markets reveals that the price volatility exhibits a strong maturity effect for three consumption commodities (crude oil, natural gas and corn) and significant seasonal variation for commodities with strong seasonality in demand (natural gas) or supply (corn). The futures price volatility for the three consumption commodities also exhibits complex dynamics in its composition among the two common factors and the contract-specific shock; volatility increases rapidly in the last month of trading, a large share of which emanates from the contract-specific shock. Consequently, the correlation between the nearby futures price and the prices of more distant contracts declines sharply as the contract approaches its maturity date. The model also reveals that the inter-temporal price linkage breaks at the end of winter for natural gas and at around September to early November for corn. The finding supports the implication of the theory of storage that the inter-temporal price link breaks when inventory clears out, which happens at the end of the winter peak-demand season for natural gas and in September right before the new harvest arrives for corn.

The conventional two-factor Gaussian model cannot replicate these complex price dynamics due to the restrictive specifications stipulated for the stochastic processes of the underlying factors and the variance of the measurement error. In particular, the commonly considered model specification forces the short-term factor to capture entirely the variation in the futures price volatility by time-to-maturity while it assigns the measurement error to capture the seasonal and cross-contract variation in the futures price volatility. The specification also stipulates that the variance attributable to the two common factors decreases exponentially with time-to-maturity at an identical rate for all contracts. Due to these restrictions, the model implies that cross-contract correlation increases monotonically as the contract approaches maturity, which is exactly opposite to the implication of the flexible model.

The composite model, allowing a flexible variance structure of the idiosyncratic error, performs reasonably well in replicating the complex price dynamics depicted by the flexible model. The results highlight that specifying a flexible variance structure of the idiosyncratic error alone can improve the ability

of conventional term-structure models to replicate the complex volatility dynamics of consumption commodities. This finding cautions against a recent trend in the development of term-structure models, which focuses exclusively on adopting more flexible stochastic processes of common stochastic factors while maintaining restrictive variance structures of the measurement error.

An incorrect portrayal of volatility and cross-contract correlation can lead the conventional term-structure model to suggest hedging strategies that are less effective than the strategy based on the flexible model. The strategy based on the composite model is also less effective than the flexible model, particularly for the two consumption goods with strong seasonality in demand or supply. These results indicate that specifying the variance of the idiosyncratic error alone cannot replicate properly the complex dynamics of the price correlation across concurrently traded contracts, which is critical for designing an effective hedging strategy and for accurate pricing of derivative contracts whose payoff depends on the realization of price spreads.

Acknowledgements

I gratefully acknowledge insightful comments from two anonymous reviewers, which helped improve the paper substantially. This research was conducted with financial support from an Australian Research Council Discovery Grant (DP1094499). The author benefitted from discussions with Aaron Smith, Felix Chan, Harry Bloch and participants of MODSIM 2011.

References

- Black, F. and Scholes, M., The pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**, 637–654.
- Cortazar, G. and Schwartz, E.S., The valuation of commodity-contingent claims. *J. Deriv.*, 1994, **1**, 27–39.
- Engle, R.F. and Kroner, K.F., Multivariate simultaneous generalized ARCH. *Econometr. Theory*, 1985, **11**, 122–150.
- Gibson, R. and Schwartz, E.S., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Hamilton, J.D., State-space models. In *Handbook of Econometrics*, vol. 4, edited by R.F. Engle and D.L. McFadden, 1994 (Elsevier: Amsterdam).
- Heath, D., Jarrow, R. and Morton, A., Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica*, 1992, **60**(1), 77–105.
- Hull, J.C., *Options, Futures, and Other Derivatives*, 5th ed., 2002 (Prentice Hall: Englewood Cliffs, NJ).
- Lautier, D., Term structure models of commodity prices: a review. *J. Altern. Invest.*, 2005, **8**(1), 42–64.
- Lucia, J. and Schwartz, E.S., Electricity prices and power derivatives: evidence from the Nordic power exchange. *Rev. Deriv. Res.*, 2002, **5**, 5–50.
- Manoliu, M. and Tompaidis, S., Energy futures prices: term structure models with Kalman filter estimation. *Appl. Math. Finance*, 2002, **9**, 21–43.
- Nielsen, M.J. and Schwartz, E.S., Theory of storage and the pricing of commodity claims. *Rev. Deriv. Res.*, 2004, **7**, 5–24.
- Routledge, B.R., Seppi, D.J. and Spatt, C.S., Equilibrium forward curves for commodities. *J. Finance*, 2000, **55**, 1297–1338.
- Schwartz, E.S., The stochastic behavior of commodity prices: implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Schwartz, E.S. and Smith, J.E., Short-term variations and long-term dynamics in commodity prices. *Mgmt Sci.*, 2000, **46**, 893–911.
- Smith, A., Partially overlapping time series: a new model for volatility dynamics for commodity futures. *J. Appl. Econometr.*, 2005, **20**, 405–422.
- Sorensen, C., Modeling seasonality in agricultural commodity futures. *J. Fut. Mkts*, 2002, **22**(5), 393–426.

- Suenaga, H. and Smith, A., Volatility dynamics and seasonality in energy prices: implications for crack-spread price risk. *Energy J.*, 2011, **32**(3), 27–58.
- Williams, J.C. and Wright, B.D., *Storage and Commodity Markets*, 1991 (Cambridge University Press: New York, NY).

Application of a TGARCH-Wavelet Neural Network to Arbitrage Trading in the Metal Futures Market in China

14.1	Introduction	280
14.2	Model Development: TGARCH-WNN Statistical Arbitrage	281
	Mean Spread with TGARCH • Wavelet Theory and Chaotic Properties of Time Series • Wavelet Neural Networks	
14.3	Sample Selection	285
14.4	Empirical Study.....	287
	Application of the TGARCH Model • Testing Chaotic Properties Using Lyapunov Exponents • Robustness Check of the TGARCH- WNN, Backpropagation Neural Network and Historical Optimal Methods	
14.5	Conclusion	297
	Acknowledgements	297
	References.....	297

Lei Cui

Ke Huang

H.J. Cai

For high-frequency statistical arbitrage, setting the proper trading threshold for each trading period is extremely important. We find that the optimal short-run trading threshold series in the Chinese metal futures market demonstrates chaotic characteristics. Therefore, we propose a new statistical arbitrage model in which a TGARCH model is applied to capture short-term price cointegration and asymmetric price spread standard deviations between futures contracts and a wavelet neural network is utilized to predict trading thresholds. Backtesting results demonstrate that the new model provides more stable and accurate trading thresholds and therefore generates more profits compared with the backpropagation neural network and historically optimal models. Potential topics for further study, including adjustments of the model to fit different market situations and the impact of transaction costs, are discussed.

Keywords: Statistical arbitrage; TGARCH; Wavelet neural network; Futures market; Commission level

JEL Classification: C, C61, C45

14.1 Introduction

Worldwide, metal futures markets play an indispensable role in both individual and corporate investment. During its liquidity boom, the metal futures market in China has been rife with fraud and price manipulation because of the inefficient market environment.* This feature makes statistical arbitrage an appropriate quantitative trading strategy in this market. Generally, a statistical arbitrage consists of three main parts (Steve *et al.* 2004): obtaining arbitrage pairs, establishing the spread series and setting trading thresholds and stopping limits.

The GARCH-based financial time series model, which was especially designed to address volatility clustering, has proven to be successful at modelling many aspects of financial time series (Bollerslev *et al.* 1992, Weber and Prokopczuk 2011). Many studies have applied the GARCH model to determine ideal trading pairs and to predict the spread standard deviation in statistical arbitrage. However, when applying the GARCH model to predict the spread standard deviation, Christian and Zakoian (2010) found that positive and negative information have a different influence on prediction, which is consistent with the asymmetric behaviour commonly exhibited by the stock market (Wang and Wang 2011). Thus, asymmetric GARCH models have been utilized to analyse the volatility of financial time series. He (2008) and Lin *et al.* (2010) applied the threshold GARCH (TGARCH) model to study asymmetric effects in the Chinese metal futures market. Their works, however, primarily study daily trading data to determine long-term equilibrium asymmetric effects that exist in the market. In our study, attention is paid to short-term asymmetric effects, which are more meaningful for high-frequency trading models.

For high-frequency statistical arbitrage, for which the holding period is just a few minutes, setting the proper trading threshold is very important. Therefore, the traditional method to set the trading threshold, which utilizes the historical optimal value and is described by Gatev *et al.* (1999), becomes unreasonable. Following the method established by Kantz (1994), we study the characteristics of the optimal trading threshold series which maximize profits and note its chaotic properties, which leads us to consider using learning algorithms, especially neural networks, to predict arbitrage trading thresholds. Neural networks have been widely applied to predict price trends in equity markets using daily prices (Creamer 2012). However, McMillan and Speight (2006) argue that for the strong microstructure relationship between informed and noise traders in futures markets, it is risky to apply neural networks alone in trading models. One of the most effective methods to solve this problem is to utilize the wavelet transform to alter the noise during the signal process (Zhang *et al.* 2001, Jaisimha *et al.* 2010, Chang *et al.* 2011). Thus, Li and Kuo (2008) suggest that wavelet neural networks (WNNs) are a promising prediction model for establishing more effective trading models.

Studies by Locke and Venkatesh (1997) argue that transaction costs are an essential factor for investments in futures market. Costs influence the market liquidity (Brennan and Subrahmanyam 1996) and bid-ask spread (Glosten and Harris 1988). For quantitative trading methods such as statistical arbitrage, transaction costs, especially the commission level, can even determine the profitability of arbitrage (Chung 1991). Thus, the impact of changes in the commission level on the cumulative yields and total commission in our model will be discussed in detail.

Previous studies concerning quantitative arbitrage models in the Chinese futures market focused on the cointegration relationships among different types of contracts. For instance, Fung *et al.* (2010) established a trading strategy based on the cointegration relationship between aluminium and copper futures price series in the Shanghai Futures Exchange (SHFE). However, such methods fail to provide accurate trading recommendations. Thus, our study combines the advanced theories of non-linear dynamics and artificial neural networks with a traditional statistical arbitrage model to create a new arbitrage strategy for metal futures markets. There are two primary contributions of our study. First, based on the sketch model proposed by Huang *et al.* (2012), we apply threshold GARCH (TGARCH) models to

* Detailed information about Chinese futures markets can be found at '<http://www.world-exchanges.org/news-views/views/chinese-futures-markets-coming-booming-decade>'.

capture short-term price cointegration and extend the model to the metal futures market in China. We also further study adjustments of the model to fit different market situations based on the TGARCH model. Second, to the best of our knowledge, our model is the first to predict trading thresholds in the Chinese metal futures market using WNNs. With deliberately designed backtesting programs that take the margin rate, margin call, slippage and commission levels into account, we use high-frequency data to demonstrate the robustness of the WNN model compared with typical predication methods, such as the backpropagation neural network (BP) and historical optimal value (HO) models, and to test the profitability of statistical arbitrage using TGARCH-WNN under different commission levels.

The remainder of the paper is organized as follows: first, an outline of the TGARCH-WNN statistical arbitrage model, which combines cointegration with the TGARCH and WNN models, is given. Then, the high-frequency data and contract pairs used in this paper are described in the sample selection section. The final part of the paper is the empirical study section, which contains four subsections. In the first subsection, discussions about the application of the TGARCH model to eliminate heteroscedasticity and an analysis of asymmetric effects under distinctive market situations are presented. In the second subsection, the Lyapunov exponent is utilized to demonstrate the chaotic properties of the trading threshold series. In the third subsection, a comparison of the prediction accuracy and market performance of the WNN, BP and HO models is made. Finally, the impact of commissions is discussed in detail.

14.2 Model Development: TGARCH-WNN Statistical Arbitrage

There are two parts to the TGARCH-WNN model. In the first part, TGARCH is applied to find the cointegration relationship between contracts and to predict the standard deviation of the price spread. In the second part, WNNs are used to predict the trading thresholds.

14.2.1 Mean Spread with TGARCH

Without loss of generality, it is assumed that there are only two contracts X and Y in the arbitrage pair and that the price series of X and Y are x and y , respectively.*

- y serves as the dependent variable and x serves as the independent variable. If y is cointegrated with x , then x and y are a pair of potential trading candidates for statistical arbitrage. Moreover, x and y will be regressed using TGARCH (1, 1) to avoid the influence of heteroscedasticity in the time series. The threshold GARCH model is applied here because positive news and negative news have distinct impacts on variance prediction. The regression functions are as follows:

mean function:

$$y_t = \beta_0 + \beta_1 x_t + u_t \quad (14.1)$$

variance function:

$$\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \gamma u_{t-1}^2 d_{t-1} + \beta \sigma_{t-1}^2, \quad (14.2)$$

where u_t are the disturbance terms in the mean function, σ_t is the conditional standard deviation series and d_t is determined by u_t . When $u_t < 0$, $d_t = 1$ and when $u_t \geq 0$, $d_t = 0$. In the mean function (14.1), positive news ($u_t \geq 0$) and negative news ($u_t < 0$) have different influences on the conditional variance: the influence level of positive news is α ; the influence level of negative news is $\alpha + \gamma$. If $\gamma > 0$, there is a leverage effect, which means the asymmetric effect will increase the volatility of variance; if $\gamma < 0$, the asymmetric effect will decrease the volatility of variance.

* When there are more than two contracts $\{y_t, x_{1t}, x_{2t}, \dots, x_{nt}\}$, we define their mean function as $y_t = \beta_0 + \beta_1 x_{1t} + \dots + \beta_n x_{nt} + u_t$

- From the above model, we can obtain residual series:

$$u_t = y_t - \beta_0 - \beta_1 x_t. \quad (14.3)$$

There is evidence that the Phillips–Perron (PP) test has more power than the augmented the Dickey–Fuller test to examine stationarity.* Thus, the PP test is applied to check whether the residual series of (14.3) is stationary. If it is stationary, it can be concluded that there is a cointegration relationship between $\{y_t\}$ and $\{x_t\}$, which means that this contract pair can be used for statistical arbitrage. Let $\{1, -\beta_1\}$ be the integrated vector and $\{u_t\}$ be the spread series which is renamed as $\{spread_t\}$.

- If the $\{spread_t\}$ passes the PP test, $\{spread_t\}$ is decentralized such that the equilibrium value of the series is zero. Such a residual series is denoted as mean spread: $mspread_t = spread_t - \text{mean}(spread_t)$

Before moving forward to WNN, the general idea of statistical arbitrage is presented here. As Figure 14.1 shows, in the process of arbitrage, the $mspread_t$ series of the contract pair (denoted as e_t) is monitored over time. Let $k_1\sigma_t$ and $k_2\sigma_t$ be the values of the trading thresholds (the dashed lines in Figure 14.1), where σ_t is the standard deviation of $mspread_t$ at time t , and can be predicted from function 14.2. Trading is triggered when $|e_t|$ deviates from the equilibrium by more than $|k_1\sigma|$ or $|k_2\sigma|$, namely $e_t > k_1\sigma$ or $e_t < -k_2\sigma$. Let x and y be the contracts defined in Equation 14.1. The detailed trading rules are:

When $e_t > k_1\sigma$, buy β contracts of X and sell one contract of Y . Close the positions when e_t decreases below zero.

When $e_t < -k_2\sigma$, sell β contracts of X and buy one contract of Y . Close the positions when e_t increases above zero.

However there are times when $mspread$ deviates far away from the average. This situation is very dangerous, so stop losses (the solid lines in Figure 14.1) should be set to avoid large money losses. Let the upper and lower stop losses be s_1 and s_2 , then positions should also be closed when $e_t > s_1$ or $e_t < -s_2$. The distribution of $mspread$ is assumed to be $N(0, \sigma_t)$ (Noortwijk *et al.* 2007). Thus let $s_1 = n\sigma_t$ and $s_2 = -n\sigma_t$ where $n\sigma_t$ is the significance value of a small probability event.

14.2.2 Wavelet Theory and Chaotic Properties of Time Series

Now, the only question left is to predict the optimal $\{k_1, k_2\}$. For simplicity, we call the $\{k_1, k_2\}$ series the trading threshold series in the following section. This paper applies WNN to predict the appropriate




FIGURE 14.1 Arbitrage process of a contract pair.

* Concrete evidence can be found in study by (Davidson and Mackinnon 1993).

trading threshold. To introduce WNN, two points must be addressed first: what is wavelet theory and why it is sensible to apply WNN to predict trading thresholds?

Wavelet theory is a new method to transform signals into simpler and lower dimensional formations. Unlike the traditional Fourier transform, which loses time information when transforming a time series, the wavelet transform accurately retains time information. Moreover, the wavelet transform is useful for eliminating noise when predicting the trading threshold value.

It is well known that neural networks are an ideal method to make predictions when time series have chaotic properties. Therefore, the chaotic properties of trading threshold series should be confirmed before using WNN for prediction.

The Lyapunov exponent is applied to prove that the trading threshold series has chaotic properties. If the Lyapunov exponents of the series have stable positive values as the embedding dimension increases, which is also defined as having positive largest Lyapunov exponents, then the series has chaotic properties. The detailed process can be found in the study by Wolf *et al.* (1985) and the concise procedure is presented as follows:

- The k series is denoted by $\{k(t), t = 1, 2, 3, \dots, M\}$.^{*} To get the average period, the k series is transformed by Fast Fourier Transform (FFT) and get: $F(u) = \sum_{i=1}^M k(i)e^{-j(u-1)f_i}$ where j is an imaginary number. In the transformation, the frequency function is: $f_i = 2\pi(i-1)/M$.
- In order to find the optimal point k^* of the phase space, we minimize the Euclidean distance $L(t')$ between $K(t')$ and $K(t)^*$ such that $\|t - t'\| > T_m$, where t' denotes the base point and satisfies $t'_n = t'_{n-1} + q$, $n = 1, 2, \dots, M - (m-1)\tau$, $t = 1, 2, \dots, t' - 1, t' + 1, \dots$
- Increase the embedding dimension m until Lyapunov exponents[‡] become stationary which is defined as Largest Lyapunov exponent. If Largest Lyapunov exponent converges to a positive number, then $\{k(t), t = 1, 2, 3, \dots\}$ has the properties of chaos.

Here two improvements have been done to meet the special character of arbitrage spread series. Firstly, historical trading thresholds $k(t)$ which maximize profits in each period make up the time series. Secondly, when calculating Lyapunov exponents, the traditional way to find the average period is unconvincing with regard to the special time series here. Instead, frequency-weighted average period is calculated by (14.4) (Ellen *et al.* 1998).

$$T = \frac{\sum_{i=2}^M \frac{F(i)}{f_i}}{\sum_{i=2}^M f_i} \quad (14.4)$$

14.2.3 Wavelet Neural Networks

In this paper, a tight WNN is applied to make a one-step-ahead forecast. This type of neural network is based on the typical BP neural network and applies the wavelet basis function as the hidden nodes transfer function. Apart from wavelet nodes, some traditional multilayers of BP neural networks are added in the hidden layers to increase the convergence rate. In the network, signals spread forward, whereas error series spread backward, as shown in Figure 14.2.

In Figure 14.2, $K_1, K_2, \dots, K_n$ are WNN's input parameters.[§] $Y_1, Y_2, \dots, Y_m$ are WNN's predictions, ω_{ij} and ω_{jk} are, respectively, wavelet nodes' input and output weight.

^{*} There is an upper k and lower k for each period and the k series here means either the upper or lower k series.

[†] $K(t)$ is the reconstructed phase space: $K(t) = \{k(t), k(t+\tau), \dots, k(t+(m-1)\tau)\}$, where m is the embedding dimension and τ is the delay time which equals the time step q .

[‡] The Lyapunov exponent can be written as $LE(m) = \frac{1}{m} \sum_{i=1}^m \frac{1}{k} \log_2 \frac{L(t_i)}{L(t_{i-1})}$.

[§] In our model input parameters are $\{k(t), t = 1, 2, 3, \dots\}$




FIGURE 14.2 Architecture of WNN.

When the input signals are k_i ($k_i \in K_i$, $i = 1, 2, \dots, n$), the wavelet functions are:

$$f(j): f_j \left[\frac{\sum_{i=1}^n \omega_{ij} x_i - \gamma_j}{\alpha_j} \right] \quad j = 1, 2, \dots, l. \quad (14.5)$$

In the function (14.5) above, $f(j)$ is the output of j 's hidden node; f_j is a wavelet basis function.

To decide the threshold in the WNN, the Mexican Hat Function is applied as mother function (Zhou and Hojjat 2003): $y(x) = \frac{2}{\sqrt{3}} \pi^{-\frac{1}{2}} (1 - x^2) e^{-\frac{x^2}{2}}$

The scales and shifts for each wavelet node are optimized by including them as parameters within the usual training algorithm as in a multilayer BP neural network. Thereby the optimal set of scales and shifts appropriate for the prediction task is determined.

Utilizing the theory above, the main process of WNN prediction is as follows:

- Initializing the WNN: Assign initial value to parameters in the network structure including learning rate, epoch size, shifts and scales.
- Data preprocessing: Determine the size of testing and training sets. In the model, optimal $\{k(t), t = 1, 2, 3 \dots\}$ series is the same series mentioned in Section 14.2.2. Denote the value of the k series as k_{it} , where $i = 1, 2$; $t = 1, 2 \dots, m$ ($i = 1$ means upper threshold and $i = 2$ means lower threshold). Thus, k series of m periods are used for training are $\{k_{11}, k_{12}, k_{13}, \dots, k_{1m}\}$. If n sample points are used to predict one point, the series should be reconstructed, as shown in Table 14.1.
- Training WNN: The input the series shown in Table 14.1 input is into the WNN to get the prediction vector $\{k'_{in+1}, k'_{in+2}, k'_{in+3}, \dots, k'_{im}\}$. Compare $\{k'_{in+1}, k'_{in+2}, k'_{in+3}, \dots, k'_{im}\}$ with $\{k_{11}, k_{12}, k_{13}, \dots, k_{1m}\}$ and calculate the value of target function $E(\theta)$:

$$E(\theta) = \frac{1}{2} \sum_{j=n+1}^m (k'_{ij} - k_{ij})^2 \quad i = 1, 2. \quad (14.6)$$

- According to the value of $E(\theta)$, revise the parameters in the WNN until $E(\theta)$ is smaller than the adaptive tolerance.
- Predict the k value and apply the well-trained WNN to predict k 's trading threshold value.

TABLE 14.1 Input and Target of WNN

Input	Target Value
$k_{i1}, k_{i2}, k_{i3}, \dots, k_{in}$	$k_{i(n+1)}$
$k_{i2}, k_{i3}, k_{i4}, \dots, k_{i(n+1)}$	$k_{i(n+2)}$
$k_{i(m-n)}, k_{i(m-n+1)}, k_{i(m-n+2)}, \dots, k_{i(m-1)}$	k_{im}

14.3 Sample Selection

The empirical study in this paper is focused on the commodities in the SHFE, which is an exchange in China that primarily trades metals. Taking liquidity and trading volume into account, only primary contracts are considered. A typical calendar spread arbitrage pair, zinc, is selected to conduct empirical testing for two practical reasons.* First, zinc's relatively high trading volume ensures that there are enough market orders for us to set up positions when the trading signal is triggered. Second, unlike deformed steel bars and copper, for which the commission level is based on a percentage of the transaction volume, the commission level of zinc is based on the number of lots in the transaction.† When the commission level is based on the number of lots, it is easier to calculate parameters such as the lower bound of the trading threshold and the total commission.‡

Generally, two contracts are chosen to construct zinc calendar spread arbitrage: the primary contract, which has the largest trading volume, and the secondary contract, which has the second largest trading volume. Because the trading of the primary contract is more active than that of the secondary contract, it is reasonable to assume that the price of the secondary contract is influenced by the price of the primary contract. Thus, the price of the primary contract serves as the independent variable, whereas the price of the secondary contract serves as the dependent variable.

Data of different frequencies are applied in each part of the empirical studies to fit some special needs. When analysing the cointegration relationships of the contract pairs using the TGARCH model, it is essential to capture the characteristics of short-term equilibrium and avoid the noise in the price series. Thus, closing price series sampled at one-minute intervals are used here. It is important to apply high-frequency data to test the performance of quantitative trading models, such as the statistical arbitrage model in this paper. Thus, tick-by-tick intraday data from SHFE are utilized in the trading simulation. We trade and quote level two data, which means that the sample series in the trading simulation are recreated at one-second intervals. Each point in the series is the volume-weighted average price (bid price, bid volume, ask price and ask volume).

To make the empirical studies more realistic, four points should be emphasized here. First, according to the standard of the SHFE, the contract size of zinc is 5 tons per lot and the margin rate is assumed to be 10%. Second, the trading hours each day are separated into three parts, period 1 (9:00–10:15), period 2 (10:20–11:30) and period 3 (13:30–15:00). A pair of trading thresholds is calculated for each period in a trading day. For example, there were 20 trading days in May 2010; thus, there are 60 points for each k series. Third, the contract pair can be utilized to perform statistical arbitrage only when the mspread series is stationary; thus, the mspread series generated using two-day data before the trading date is used to perform the PP test. For example, if we decide to do the arbitrage on 2011/11/21, then we utilize the data from 2011/11/19 to 2011/11/20 to conduct the PP test. If the series is proven to be stable, then we conclude that the cointegration relationship will still hold on 2011/11/21. Fourth, to demonstrate the effectiveness of TGARCH-WNN, the model should be examined under distinct market circumstances. Therefore, three typical trading periods are selected to conduct the empirical study. The characteristics of these periods are described in [Table 14.2](#) and [Figure 14.3](#).


* The TGARCH-WNN model can be applied to all primary contracts in the SHFE.

† In the SHFE, the commission level of deformed steel bars is three ten thousandths of the turnover, whereas the commission level of copper is one ten-thousandth of the turnover.

‡ A detailed discussion of this issue is in Section 14.4.4.

TABLE 14.2 Market Conditions of Three Typical Trading Periods

Trading Period	Contract Pair	Market Conditions
2010/07/20–2010/07/23	zn09&zn10	Both contracts are falling continuously
2010/04/27–2010/04/30	zn06&zn07	Both contracts are rising continuously
2011/11/21–2011/11/24	zn01&zn02	Both contract prices are stable

**FIGURE 14.3** Sample price movement of zinc calendar spread arbitrage.*(Continued)*




FIGURE 14.3 (Continued) Sample price movement of zinc calendar spread arbitrage.

The samples selected here are all cointegrated pairs of stationary price series that passed the PP test. Like other statistical arbitrages, the trading frequency of the TGARCH-WNN arbitrage model is high and all positions are closed within the day; thus, long-term market situations have little influence on the arbitrage. That is the reason why only short-term market trends are considered here. Therefore, each market situation contains four days of tick-by-tick data.

14.4 Empirical Study

The empirical study is separated into four parts. In the first part, a discussion of the application of the TGARCH model to eliminate heteroscedasticity and an analysis of asymmetric effects under distinct market situations are presented. In the second part, the Lyapunov exponent is utilized to prove the chaotic properties of the k series. In the third part, a comparison of the prediction accuracy and market performance of the WNN, BP and HO models is made. Finally, the impact of commission level is discussed in detail.

14.4.1 Application of the TGARCH Model

As mentioned in [Section 14.2.1](#), to eliminate heteroscedasticity in the regression and to take asymmetric effects into consideration when predicting the $mspread$ series' standard deviation,* the TGARCH regression model is applied. Here statistical tests are conducted to demonstrate the effectiveness of TGARCH at removing heteroscedasticity and present detailed characteristics of asymmetric effects in different market situations. Three typical market situations are considered.† Consistent with the points mentioned in [Section 14.3](#), two-day minute data are selected to conduct statistical tests.

[Table 14.3](#) presents the ARCH-LM test results of the TGARCH regression and OLS regression of contract pairs in different market situations. The null hypothesis of the ARCH-LM test is that there is no

* As mentioned in Equation (14.2).

† As shown in [Table 14.2](#) and [Figure 14.3](#).

heteroscedasticity in the regression. From Table 14.3, we can see that the p values of the OLS regression are zero, which means that heteroscedasticity is apparent in all market situations. In contrast, the tests results of the TGARCH regression have significantly positive p values, which means TGARCH regression eliminates heteroscedasticity effectively.

It is important to understand the properties of asymmetric effects under different market circumstances when predicting the standard deviation of price spread. Thus, we examine the significance of the $u_{t-1}^2 d_{t-1}$ term in the function (14.2) and the sign its coefficient.

From Table 14.4, the p value of $u_{t-1}^2 d_{t-1}$'s coefficient in the periods 2010/04/26–2010/04/27 and 2010/07/21–2010/07/22 is less than 0.05; thus there are asymmetric effects when the price of both contracts increases or decreases unidirectionally. However, the coefficient of $u_{t-1}^2 d_{t-1}$ in the two periods has different signs. In the period 2010/04/26–2010/04/27, when the contract prices demonstrate a downward trend, $u_{t-1}^2 d_{t-1}$'s coefficient has a positive value which means that the asymmetric effect increases the variance of the price spread, whereas in the period 2010/07/21–2010/07/22, when the contract prices increase unidirectionally, the asymmetric effect decreases the variance of the price spread because of the negative value of $u_{t-1}^2 d_{t-1}$'s coefficient. This result is reasonable because the unidirectional change of the contract price implies that the powers of positive news and negative news are different. When price increases unidirectionally, positive news dominates the market; when price decreases unidirectionally, negative news dominates the market. Thus, the influence of the news on the variance of the price spread is distinct under these circumstances.

However, in the period 2011/11/21–2011/11/22, when the zinc price is stable, the p value of $u_{t-1}^2 d_{t-1}$'s coefficient is much greater than 0.05 which means the asymmetric effect is not apparent in this period. The reason is because positive news and negative news have the same power in the market. Then, Equation (14.2) can be simplified to

$$\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2, \quad (14.7)$$

The result of the asymmetric effect analysis above is critical, because it helps to adjust the trading model with the TGARCH-WNN trading model to fit different market situations: when the contract

TABLE 14.3 ARCH-LM Test (lags:2) Result Comparison

Sample information		Test result of TGARCH		Test result of OLS	
Contract pair	Trading time	F -statistics	Prob. $F(2450)$	F -statistics	Prob. $F(2450)$
zn06&zn07	2010/04/26–2010/04/27	0.103336	0.9018	20.73649	0.0000
zn09&zn10	2010/07/21–2010/07/22	0.148928	0.8617	19.62015	0.0000
zn01&zn02	2011/11/21–2011/11/22	0.242065	0.7851	100.4303	0.0000

TABLE 14.4 Regression Results

Parameters ^a	Coefficient	p -value	Coefficient	p -value	Coefficient	p -value
	2010/04/26–2010/04/27		2010/07/21–2010/07/22		2011/11/21–2011/11/22	
Mean function:						
x_t	1.027	0.0000	1.028	0.000	1.173	0.0000
Constant	−305.148	0.0000	−378.483	0.0000	−2621.516	0.0000
TGARCH residuals:						
u_{t-1}^2	−0.052	0.0124	0.501	0.0000	0.210	0.0025
$u_{t-1}^2 d_{t-1}$	0.109	0.0000	−0.506	0.0000	−0.057	0.4759
σ_{t-1}^2	0.518	0.0000	−0.011	0.0000	0.629	0.0000
Constant	4.060	0.0000	26.558	0.0000	6.540	0.0212

^a The parameters in the table are according to function (14.1) and function (14.2)

price is changing unidirectionally, (14.2) is applied to predict *mspread*'s standard deviation, whereas when the contract price is stable, (14.7) is utilized to forecast *mspread*'s standard deviation.

14.4.2 Testing Chaotic Properties Using Lyapunov Exponents

As mentioned in Section 14.2.2, to demonstrate that it is rational to predict optimal trading threshold series using WNN, the chaotic properties of the sample series are tested by calculating the Lyapunov exponents of each series. Because the chaotic property is a long-term property of the trading threshold series, instead of applying samples from three typical situations, two-month data that contain different market circumstances are used to establish *k* series and calculate the Lyapunov exponents.* Specifically, we apply tick-by-tick data from March 1, 2010 to April 30, 2010 to generate the optimal trading threshold series in the test.

Figure 14.4 and Table 14.5 show the details of the test. From Figure 14.4 and Table 14.5, the testing series has stable positive largest Lyapunov exponents, which demonstrates the properties of chaos in such periods. As mentioned above, two months of data contain sufficient different market conditions. Thus, we can reason that the three typical market conditions in Section 14.3 will also have chaotic properties, which is the premise for the following forecast of the optimal trading threshold series.




FIGURE 14.4 Lyapunov exponents.

TABLE 14.5 Lyapunov Exponents of Two Contract Pairs

Embedding Dimension ^a	LE	Embedding Dimension	LE
kl of LE:			
32	0.0071	36	0.0054
33	0.0067	37	0.0054
34	0.0055	38	0.0055
35	0.0055	39	0.0057
k2 of LE:			
32	0.00029	36	0.00016
33	0.00044	37	0.00012
34	0.00018	38	0.00023
35	0.00022	39	0.00020

^aBased on our data test, only when embedding dimension increases to above 30 do Lyapunov exponents tend towards stability. We thus display the beginning embedding dimension from above 30.

* As mentioned in Section 14.3, only contracts with the largest trading volumes are used to conduct arbitrage. Because the primary contracts of zinc futures shift monthly, the contracts used to calculate the *k* series are different each month. In addition, the short-run market conditions alter several times within a month according to our backtesting. Thus, based on the above two points, it is reasoned that two months of data contain enough different market conditions.

14.4.3 Robustness Check of the TGARCH-WNN, Backpropagation Neural Network and Historical Optimal Methods

In traditional statistical arbitrage models, trading thresholds are trading signals that maximize the cumulative yield historically in the latest trading term. This method is called the historical optimal (HO) method. The backpropagation (BP) neural network, as a major artificial intelligence method, has been widely used for financial time series prediction. Thus, to demonstrate the robustness of the WNN model, comparisons among one-step-ahead forecasts specified by WNN, BP and HO in terms of prediction accuracy and return-rate performance under different market situations are conducted in back-testing. To construct optimal architecture networks for the WNN and the BP neural network, trading threshold series generated from four-months data (from December 1, 2009 to March 31, 2010) are utilized for training and trading threshold series obtained from the three periods described in [Section 14.3](#) are used for out-of-sample performance evaluation. The optimal parameters of WNN are presented in Table 14.6. For the BP neural network, the same statistical tests are conducted; the result is almost the same as for BP. The only difference is that the BP network contains two hidden layers, with 13 and 11 nodes, respectively.

14.4.3.1 Comparison of Prediction Accuracy

The prediction series generated by the three models are compared with the target series, which maximizes the cumulative yield in each period. The mean absolute positive error* is utilized as the standard of prediction accuracy. Table 14.7 shows the results of the comparison.

From Table 14.7, the predictions from WNN have smaller absolute positive errors than those from BP and HO in all three periods. To see how accurate the prediction of the trading thresholds is for the three models, some typical prediction points are shown in [Figure 14.5](#).[†] It is obvious that the predictions of the trading thresholds in the HO model deviate seriously from those of the other two models. Thus, only the accurate predictions of the BP and WNN models are discussed in detail here. From Figure 14.5,

TABLE 14.6 Architecture Network of WNN

Input Nodes	Output Nodes	Node Parameters			
		Weight Study Rate	Parameter Study Rate	Momentum Study Rate	Iteration
6	1	0.01	0.001	0.7	5000
Hidden layer					
Hidden layer		Wavelet		BP neural network	
Number of layer		1		2	
Number of hidden nodes		8		8–10	

TABLE 14.7 Accuracy Test Comparison of Three Models

Testing Period	K	WNN (%)	BP Neural Network (%)	Historical Optimal (%)
2010/04/27–2010/04/30	k1	19.72	24.71	30.14
	k2	16.98	20.37	23.49
	k1	16.68	18.12	22.71
2010/07/20–2010/07/23	k2	17.49	18.37	22.47
	k1	17.90	21.56	23.46
	k2	17.46	22.24	18.16

* As mentioned in function (14.6).

[†] For a clearer comparison, only the data from 2010/7/20 to 2010/7/23 are shown in the figure.




FIGURE 14.5 Prediction accuracy result of contract pair zn09 & zn10.

TABLE 14.8 Backtesting Results for Three Trading Models

Trading Models	Number of Trades	Trading Frequency (per min)	Winning Rate ^a (%)	Cumulative Yield ^b (%)
2010/04/27–2010/04/30:				
WNN	485	0.539	89.42	6.14
BP	546	0.607	88.64	5.35
HO	419	0.466	86.81	4.86
2010/07/20–2010/07/23:				
WNN	789	0.877	90.11	54.580
BP	675	0.750	87.70	50.400
HO	617	0.686	86.22	48.155
2011/11/21–2011/11/24:				
WNN	1259	1.410	90.39	105.330
BP	1299	1.440	90.07	89.005
HO	1004	1.12	87.75	78.279

^a *Winning rate* = $(m/n) \times 100$ where m is number of trade that yields positive profit (*off set* – *transaction fee* > 0) and n is total number of trade.

^b The initial margin in the test is 100000 yuan.

although sometimes the prediction of BP is closer to the optimal value than is the prediction from WNN, the BP prediction is seriously affected by the noise in the mspread series. Thus, its prediction value series is more volatile, whereas the prediction value series from WNN is more stable and less affected by noise in mspread series because of the wavelet transform.

However, the mean absolute positive error does not necessarily reflect the information value of the forecast because it is a measure of fit rather than a measure of profitable signalling. Therefore, we continue to test the market performance of these models.

14.4.3.2 Comparison of Rate of Return Under Different Circumstances

Robustness checks of the rate of return of WNN, BP and HO are performed under different market conditions. The commission level applied here is the average commission level asked by futures brokers, which is 1.25 times that specified by SHFE, 10 yuan per lot. Three essential assumptions should be emphasized here. First, to make the test more realistic, all positions are closed within the trading day to avoid margin calls. Second, all the money earned is put into the arbitrage process again. Third, the slippage loss per trade is 2 points.

From Table 14.3 and Figure 14.6, the TGARCH-WNN model performs much better than the other two trading models in the three typical periods in two aspects: it gains the highest profit and it has the highest winning rate.

Apart from the comparison result, other important phenomena should be emphasized. First, comparing the cumulative yield column and trading frequency column for the three periods, which are shown in Table 14.3, all trading models gain much more money and have a higher trading frequency in period 2011/11/21–2011/11/24 than in the periods 2010/04/27–2010/04/30 and 2010/07/20–2010/07/23. This means that although TGARCH-WNN has positive return in all three periods, the time when both contract prices are stable is the most ideal situation to conduct statistical arbitrage using the TGARCH-WNN model. Second, special attention should be paid to the testing result for the period 2010/04/27–2010/04/30. From the first picture in Figure 14.6, there is a sudden decrease in the cumulative yield. This is because there is a large jump in the price at approximately 10:00 a.m. on April 30. The abrupt change of price makes the system lose approximately 30 000 yuan in a single trade. Thus, applying statistical arbitrage alone can be quite risky, and some risk management strategies should be combined with the arbitrage system for real trades.






FIGURE 14.6 Cumulative yields of different prediction methods under three typical market situations.

14.3.3.4 The Impact of Commission

In this section, an essential factor in statistical arbitrage, the commission level, is discussed. Because the trading frequency of the arbitrage is high, the level of commission has a great influence on the profit earned. To make a profit, the yield of the arbitrage must be greater than the commission charged per trade. Thus, when calculating the trading threshold, the unit commission level determines the lower bound of the trading threshold, which in turn affects the result of the WNN. However, the exact impact of the commission level is still obscure. To give a clear view of the impact, the zinc calendar spread arbitrage's performance under different commission levels is tested in the typical situations mentioned in [Section 14.3](#). In the test, the commission level starts from the minimum level specified by SHFE (8 yuan/lot) and increases by 1 yuan each time to 14 yuan per lot. In the backtesting, special attention is paid to the changes in cumulative yield, number of trades, trading frequency and total commission, which are shown in Table 14.9. The backtesting results demonstrate some important phenomena and some essential phenomena deserve our special attention.

First, from [Figure 14.7](#), the impact of the commission change is distinct for the three periods. In the periods 2010/4/27–2010/4/30 and 2010/7/20–2010/7/23, the cumulative yield decreases linearly by approximately 50% each time when the commission level increases by one yuan and drops below zero when the commission is approximately 13 yuan per lot. In the period 2011/11/21–2011/11/24, the cumulative yield decreases exponentially with a declining decreasing rate. In general, however, statistical arbitrage with TGARCH-WNN is very sensitive to changes in the commission level.

TABLE 14.9 Back Testing Result of Different Commission Levels in Three Typical Period

Commission Level (yuan Per Lot)	Cumulative Yield (%)	Number of Trade	Trading Frequency (%)	Total Commission (Yuan)
2010/04/27–2010/04/30:				
8	28.236	440	0.489	47032
9	14.560	379	0.421	43983
10	6.410	393	0.437	46890
11	−2.857	371	0.412	47201
12	−3.292	376	0.417	51396
13	−10.570	310	0.344	43160
14	−20.842	281	0.312	40320
2010/07/20–2010/07/23:				
8	81.758	721	0.801	108846
9	61.593	576	0.639	80316
10	40.380	550	0.613	78460
11	26.604	421	0.468	63723
12	8.782	416	0.462	63084
13	−1.810	351	0.390	54080
14	−9.961	268	0.297	42868
2011/11/21–2011/11/24:				
8	232.295	1168	1.298	221440
9	123.742	954	1.06	160704
10	66.445	738	0.820	115640
11	28.380	618	0.687	96910
12	5.586	505	0.561	76032
13	−3.267	384	0.427	57746
14	−2.987	273	0.303	44156




FIGURE 14.7 Cumulative yield under different commission levels in typical periods.

TABLE 14.10 Commission Level Charged by Futures Brokers

Futures Brokers	Commission Charge for Zinc (yuan per lot)
China International Futures Co., Ltd.	10
Wonder Futures	12
Baocheng Futures Co., Ltd.	10
Changjiang Futures	9
Mailyard Futures	12

Second, from discussion with futures brokers, we learned that when the trading frequency is approximately 0.5 times per minute, they generally charge less than 1.5 times the base commission level (12 yuan per lot for zinc), and the lowest commission level charged can even be 1.125 times the base commission level (9 yuan per lot for zinc). Example commissions charged by mainstream futures brokers are shown in Table 14.10. For such commission levels, statistical arbitrage with TGARCH-WNN is still profitable.

Third, from Figure 14.8, the trading frequency of in the three periods decreases linearly to approximately the same level. A mathematical reason may be provided for this phenomenon as follows:

Without loss of generality, let j be the commission level and k_{jt} be the $k1$ value under commission level j in period t (For $k2$, the notation is the same.). Let σ_t be the conditional standard deviation of $mspread$ at time t . Then the trading threshold and conditional standard deviation at t are $k_{jt}\sigma_t$ and $n\sigma_t$ respectively.* According to the trading strategy mentioned in Section 14.2.1, when $mspread_t$ equals to the upper trading threshold at time t , namely $y_t - \beta_0 - \beta_1 x_t - \mu = k_{jt}\sigma_t$, we should buy one unit of contract Y and sell β_1 units of contract X . The position should be closed when $mspread_t$ equals 0 at time t' . According to functions 14.1 and 14.2, the revenue of this trade is $y_t - y_{t'} + \beta(x_{t'} - x_t) = k_{jt}\sigma_{t'}$. However, this trade also generates a commission fee of $(1 + \beta_1)j$. Thus, to make a positive profit, $k_{jt}\sigma_t$ should be greater than $(1 + \beta_1)j$. Furthermore, it is explicitly required that the trading threshold value should be less than the stop-loss value, as discussed in Section 14.2.1. Therefore, the feasible zone of k_{jt} is $U_t = \left\{ k_{jt} \mid \frac{(1 + \beta_1)j}{\sigma_t} < k_{jt} < n \right\}$.

For commission levels i and j , if $j > i$, then $\sigma_t k_{jt} \geq \sigma_t k_{it}$ and it is harder for the $mspread$ series to hit $\sigma_t k_j$ than to hit $\sigma_t k_i$ to set the positions. Thus, the trading frequency decreases when the commission level increases. To see why the trading frequency decreases to same level, it should be noticed that the augmentation of the commission level narrows the feasible zones U_1 , U_2 and U_3 to approximately the same level. Thus, k_{j1} , k_{j2} and k_{j3} are likely to obtain approximately the same value. Moreover, the distributions of the same contract pair's $mspread$ in different periods are similar. Thus, the $mspread$ series will have nearly the same probability of hitting the trading threshold lines, which implies that the is identical, as shown in Figure 14.8.

Fourthly, it seems contradictory to perceive that the total commission generated by arbitrage decreases as the commission level increases, as shown in the total commission columns of Table 14.9. However, apart from the commission level, there are two other factors that also affect the total commission: the number of trades and the commission per trade. Their changes account for this phenomenon. To understand the phenomenon mathematically, under commission level j , let the total commission be T_j , the number of trades be N_j , the commission per trade be C_j , the real-time equity be E_j , the real-time initial margin be M_j , the real-time price be P_j and the number of contracts purchased and sold be I_j , respectively. It is assumed that all the money earned is put into the arbitrage process again; thus, $E_j = M_j$ whenever positions are set. As mentioned in Section 14.3, the margin rate is 10%; then the expression of the total commission is as follows:†

$$T_j = N_j \times C_j = N_j \times (jI_j) = N_j \times \left(j \frac{M_j}{0.1P_j} \right) \quad (14.8)$$

For commission level i and j , if $i < j$, then the arbitrage trade can earn more money with the lower commission level i each time. As time goes by, $M_i > M_j$, which in turn leads to $C_i > C_j$. It can also be observed from the number of trade columns in Table 14.9 that $N_i > N_j$. Thus, according to (14.8), higher C and N values will most likely lead to higher total commission T .

* We obtain this estimated value from the assumption that $mspread$ is normally distributed with $(0, \sigma_t)$.

† For simplicity, the strategy of how many contracts to buy and sell specifically is ignored here.




FIGURE 14.8 Trading frequency of different periods with distinctive commission level.

14.5 Conclusion

In this paper, a new statistical arbitrage model based on TGARCH-WNN is described. The empirical study focuses on the Chinese metal futures market. It is shown that the regression of futures price series is negatively impacted by heteroscedasticity. While analysing asymmetric effects under different market situations, we find that asymmetric effects are apparent when the price is trending, whereas it is faint when the contract price is stable. We apply different algorithms according to the distinct characteristics of each typical period. The robustness check in the empirical study demonstrates that when using the wavelet transform to eliminate noise in the series, the predictions made by WNN are not only more accurate and stable than those from BP and HO but also generate more profits when applied in the arbitrage process. From the commission impact test, it is shown that the TGARCH-WNN statistical arbitrage model is very sensitive to changes in the commission level. However, it is still profitable for the average commission level charged by futures brokers.

Furthermore, with highly profitable opportunities for statistical arbitrage, a statistical arbitrage model using WNN could have many potential applications and is of great commercial value. Both investors and futures companies can take advantage of this model to gain stable and considerable profit.

Finally, it is extremely significant to note that although the empirical study demonstrates that statistical arbitrage using TGARCH-WNN is very profitable, it is also quite risky. The trading model alone cannot address sudden price changes very well, which can thus cause severe losses. Therefore, to put this model into practice, a risk management system should be constructed along with the trading model.

Acknowledgements

This research was supported by Natural Science Foundation of China (11171247). We would like to thank the editor and Elizabeth Dawes, as well as the anonymous referees for their insightful comments. We would like to acknowledge the help received from the Shanghai Futures Exchange and CSMAR Solution in providing the data. The authors thanks Rui Cao and Xiaoyong Cui from Central University of Finance and Economics, Zhuo Huang from CCER in Peking University, Yulei Luo from Hong Kong University and Wei Zou from IAS in Wuhan University for many helpful comments and suggestions.

References

- Bollerslev, T., Chou, R. and Kroner, K., ARCH modelling in finance. *J. Econom.*, 1992, **52**, 5–59.
- Brennan, M.J. and Subrahmanyam, A., Market microstructure and asset pricing: On the compensation for illiquidity in stock returns. *J. Financ. Econ.*, 1996, **41**, 441–464.

- Chang, P.C., Liao, T.W., Lin, J.J. and Fan, C.Y., A dynamic threshold decision system for stock trading signal detection. *Appl. Soft Comput.*, 2011, **11**, 3998–4010.
- Christian, F. and Zakoian, J.M., *GARCH Models: Structure, Statistical Inference and Financial Applications*, 2010 (John Wiley & Sons: Hoboken).
- Chung, Y.P., A transaction data test of stock index futures market efficiency and index arbitrage profitability. *J. Finance*, 1991, **46**, 1791–1809.
- Creamer, G., Model calibration and automated trading agent for Euro futures. *Quant. Finance*, 2012, **12**, 531–545.
- Davidson, R. and Mackinnon, J.G., *Estimation and Inference in Econometrics*, 1993 (Oxford University Press: New York).
- Ellen, E.R., Norman, A.A. and Jonathan, D.B., Simplified frequency content estimates of earthquake ground motions. *J. Geotech. Geoenviron. Eng.*, 1998, **124**, 150–159.
- Fung, H.G., Liu, Q.F. and Tse, Y., The information flow and market efficiency between the US and Chinese aluminum and copper futures market. *J. Futures Market*, 2010, **30**, 1192–1209.
- Gatev, E., Goetzmann, W.N. and Rouwenhorst, K.G., Pairs Trading: Performance of a relative-value arbitrage rule. Working Paper, 1999.
- Glosten, L.R. and Harris, L.E., Estimating the components of the bid-ask spread. *J. Financ. Econ.*, 1988, **21**, 123–142.
- He, Q., Value at risk for repo interest rate based on TGARCH. *Appl. J. Financ. Econ*, Proceedings of 2008 IEEE International Conference on Service Operations and Logistics, and Informatics. IEEE/SOLI 2008, **1**, 425–428.
- Huang, K., Cui, L. and Cai, H.J., A trading model in Chinese commodity futures market using GARCH-wavelet neural networks. In *Proceedings – 2012 3rd International Conference on E-Business and E-Government*, pp. 4925–4929, 2012.
- Jaisimha, M., Imon, P. and Oleg, S., Wavelet decomposition for intra-day volume dynamics. *Quant. Finance*, 2010, **10**, 917–930.
- Kantz, H., A robust method to estimate the maximal Lyapunov exponent of a time series. *Phys. Lett. A*, 1994, **185**, 77–87.
- Li, S.T. and Kuo, S.C., Knowledge discovery in financial investment for forecasting and trading strategy through wavelet-based SOM networks. *Expert Syst. Appl.*, 2008, **34**, 935–951.
- Lin, N., Lu, H. and Zheng, Q., An empirical study on the existence of bubble in Chinese stock market: Based on TGARCH model. In *Proceedings – 2010 2nd IEEE International Conference on Information and Financial, Engineering*, pp. 87–90, 2010.
- Locke, P.R. and Venkatesh, P.C., Futures market transaction costs. *J. Futures Markets*, 1997, **17**, 229–245.
- McMillan, D.G. and Speight, E.H., Nonlinear dynamics and competing behavioral interpretations: Evidence from intra-day FTSE-100 index and futures data. *J. Futures Market*, 2006, **26**, 343–368.
- Noortwijk, J.M., Weide, J.A.M., Kallen, M.J. and Pandey, M.D., Gamma processes and peaks-over-threshold distributions for time-dependent reliability. *Reliab. Eng. Syst. Saf.*, 2007, **92**, 1651–1658.
- Steve, H., Robert, J., Melvyn, T. and Mitch, W., Testing market efficiency using statistical arbitrage with applications to momentum and value models. *J. Financ. Econ.*, 2004, **73**, 525–565.
- Wang, P. and Wang, P., Asymmetry in return reversals or asymmetry in volatilities?—New evidence from new markets. *Quant. Finance*, 2011, **11**, 271–285.
- Weber, M. and Prokopczuk, M., American option valuation: Implied calibration of Garch pricing models. *J. Futures Markets*, 2011, **31**, 971–994.
- Wolf, A., Jack, B.S., Harry, L.S. and John, A.V., Determining Lyapunov exponents from a time series. *Physica*, 1985, **16D**, 285–317.
- Zhang, B.L., Coggins, R., Jabri, M.A., Dersch, D. and Flower, B., Multiresolution forecasting for futures trading using wavelet decompositions. *IEEE Trans. Neural Networks*, 2001, **12**, 765–775.
- Zhou, Z.Q. and Hojjat, A., Time-frequency signal analysis of earthquake records using Mexican hat wavelets. *Gen. Introductory Civil Eng. Constr.*, 2003, **18**, 379–389.

III

Commodity Prices and Financial Markets

- 15 **Short-Horizon Return Predictability and Oil Prices** *Jaime Casassus and Freddy Higuera* 301
Introduction • Oil Price, the Business Cycle and Stock Returns • Short-Horizon Predictability of Stock Returns • Out-of-Sample Predictability of Stock Returns • Long-Horizon Predictability of Stock Returns • Implications for the Cross Section of Expected Returns • Conclusions
- 16 **Time-frequency Analysis of Crude Oil and S&P500 Futures Contracts** *Joseph McCarthy and Alexei G. Orlov* 337
Introduction • Literature Review • Data • Methodology • Empirical Results • Conclusions
- 17 **Sectoral Stock Return Sensitivity to Oil Price Changes: A Double-Threshold FIGARCH Model** *Elyas Elyasiani, Iqbal Mansur, and Babatunde Odusami* 359
Introduction • Data and Methodology • Empirical Findings • Conclusion
- 18 **Short-Term and Long-Term Dependencies of the S&P 500 Index and Commodity Prices** *Michael Graham, Jarno Kiviaho, and Jussi Nikkinen* 387
Introduction • Brief Literature Review • Measuring Stock and Commodities Market Dependencies Using Wavelet Squared Coherency • Data • Empirical Analysis of Comovements • Conclusion
- 19 **Portfolio Selection with Commodities Under Conditional Copulas and Skew Preferences** *Carlos González-Pedraz, Manuel Moreno, and Juan Ignacio Peña* 403
Introduction • Portfolio Choice with Commodity Futures and Skewness • Multivariate Conditional Copula with Asymmetry • Empirical Application • Conclusion
- 20 **Strategic Commodity Allocation** *Pierre Six* 433
Introduction • The Optimal Strategies Design • Empirical Assessment of the Strategies • Concluding Remarks
- 21 **Long-Short Versus Long-Only Commodity Funds** *John M. Mulvey* 459
Introduction • Sources of Alpha in Commodity Markets • Temporal Performance Measures • A Relative-Value Commodity Approach • Conclusions

22	Commodity Markets Through the Business Cycle <i>Julien Chevallier, Mathieu Gatumel, and Florian Ielpo</i>	469
	Introduction • The Reaction of Commodity Markets to Economic News • Economic Regimes and Commodity Markets as an Asset Class • Conclusion	
23	The Dynamics of Commodity Prices <i>Chris Brooks and Marcel Prokopczuk</i>	501
	Introduction • Models and Estimation • Data and Empirical Results • Economic Implications • Conclusions	
24	A Hybrid Commodity and Interest Rate Market Model <i>Kay F. Pilz and Erik Schlögl</i>	523
	Introduction • The Commodity LIBOR Market Model • Calibration With Time-Dependent Volatilities • Futures/Forward Relation and Convexity Correction • Merging Interest Rate and Commodity Calibrations • Real Data Example • Pricing Spread Options • Conclusion	

Short-Horizon Return Predictability and Oil Prices

15.1	Introduction	302
15.2	Oil Price, the Business Cycle and Stock Returns.....	303
	Oil Price and the Macroeconomy • Oil Price and the Financial Market • Stock Returns and the Business Cycle • Measuring Oil Price Shocks • Oil Price, the Business Cycle and Stock Returns	
15.3	Short-Horizon Predictability of Stock Returns	309
15.4	Out-of-Sample Predictability of Stock Returns	313
15.5	Long-Horizon Predictability of Stock Returns	316
15.6	Implications for the Cross Section of Expected Returns.....	318
15.7	Conclusions.....	329
	Acknowledgements	330
	References.....	330
	Appendix A: Tests for Out-of-Sample Predictability	333

Jaime Casassus

Freddy Higuera

This paper shows that oil price changes, measured as short-term futures returns, are a strong predictor of excess stock returns at short horizons. Ours is a leading variable for the business cycle and exhibits low persistence which avoids the fictitious long-horizon predictability associated with other predictors used in the literature. We compare our variable with the most popular predictors in a sample period that includes the recent financial crisis. Our results suggest that oil price changes are the only variable with forecasting power for stock returns. This significant predictive ability is robust against the inclusion of other variables and out-of-sample tests. We also study the cross-section of expected stock returns in a conditional CAPM framework based on oil price shocks. Our model displays high statistical significance and a better fit than all the conditional and unconditional models considered, including the Fama–French three-factor model. From a practical perspective, ours is a high-frequency, observable variable that has the advantage of being readily available to market-timing investors.

Keywords: Commodity markets; Commodity prices; Financial econometrics; Forecasting applications; Empirical asset pricing

JEL Classification: E3, E4, E32, E44, G1, G12, Q4, Q43

15.1 Introduction

The predictability of stock returns is a controversial topic. Until recently, the prevailing view was that returns could be predicted by the business cycle at long horizons (Cochrane 2005) and that this evidence was significantly stronger than for short horizons. But in recent years, several studies have questioned the existence of such return predictability. For example, Boudoukh *et al.* (2008) show that the dominant findings in this literature are solely the consequence of the high persistence of the predictors. Also recently, Welch and Goyal (2008), who compare the out-of-sample predictive performance of a large number of popular predictors with the prevailing average excess stock return, find that none of these variables predicts stock returns at short horizons better than the historical average return. The current challenge is to propose a variable that predicts equity returns at short horizons and is robust against the new tests suggested in the literature.*

This paper shows that unexpected changes in oil prices are a significant predictor of excess stock market returns at short horizons. We measure unexpected oil price shocks by short-term futures returns on crude oil contracts. Our sample period is 1983Q2–2009Q4 and is restricted by the existence of crude oil futures prices. Based on in-sample tests and on the macroeconomic literature (i.e. Hamilton and Herrera (2004)), we use four lags of this variable for the predictive regressions. Our predictive variable has deep macroeconomic roots and allows us to connect the short-horizon predictability of equity returns with the business cycle. Indeed, the existence of negative Granger-causality of oil price shocks on both equity returns and production growth in a trivariate VAR confirms that oil price shocks are a leading variable and are countercyclical. This evidence suggests that increases in oil prices precede recessions and declines in excess stock returns.

We find that, at horizons of one to three quarters, the oil price shocks exhibit predictive performance which is both statistically and economically significant. Obtaining significant results in a sample period after the oil crisis in the 1970s is not an easy task since most variables lose their forecasting power within this period (Welch and Goyal 2008). In terms of global performance, our variable is better than all other variables considered (consumption–wealth ratio, price–dividend ratio, output gap and risk-free rate) with an $\bar{R}^2$ of 6%. This meaningful in-sample result was also detected in out-of-sample tests. Our variable exhibited the best out-of-sample R^2 , which was close to 1.2% at one quarter. For longer time horizons, however, no variable showed a significant predictive performance.

Furthermore, oil price shocks have other virtues such as no persistence (they do not produce the pattern reported by Boudoukh *et al.* (2008)), they are directly observable by futures changes (unlike variables such as consumption–wealth ratio and product gap, which must be estimated), they have no correlation with the predictive regression's disturbances (do not generate the bias analysed by Stambaugh (1999)) and they are a high-frequency variable available at no cost. All of these characteristics are valued in the practice of portfolio management. To our knowledge, these results position oil price shocks as the best short-term forecasting variable today.

Our study is motivated by two lines of research: the relationship between GDP and oil prices and the linkage between equity returns and oil price shocks. Because stock return predictability has been detected at business-cycle frequencies, the first strand of literature justifies the relationship between oil prices and the macroeconomy. Intuitively, there are good reasons for believing that oil price leads the business cycle, as nine of ten recessions in the United States since World War II have been preceded by an increase in oil prices (Hamilton 2008). This has not gone unnoticed by economists and has generated a substantial amount of research, particularly given the fact that oil expenditure represents only 4% of GDP. This literature motivates the main question we address in this paper: Given that oil price shocks precede changes in GDP, do they also have some predictive power for stock market returns? The second line of related research, although less voluminous than the first, provides empirical evidence that past oil shocks have an impact on future equity returns. The following section studies these topics in more detail.

* In the words of Welch and Goyal (2008) the challenge is still open: "... our article suggests only that the profession has yet to find some variable that has meaningful and robust empirical equity premium forecasting power ...".

Our paper is also related to recent studies of short-term predictability. For example, Ang and Bekaert (2007) report evidence of predictability, at horizons up to one year, using both the dividend–price ratio and the interest rate. Campbell and Thompson (2008), using a much longer sample, find that the out-of-sample forecasting power of several variables improves significantly when certain restrictions are imposed, and that trading on these predictors can lead to significant welfare benefits when compared with trading on the historical average return. More recently, Cooper and Priestley (2009) propose the output gap as a new forecasting variable for stock market returns. This variable is robust against the tests of Boudoukh *et al.* (2008) and Welch and Goyal (2008). Finally, Bakshi *et al.* (2011) find that the Baltic Dry Index, a shipping activity variable that is tied to the business cycle, has predictive ability for a range of stock markets.

While our purpose is to study the effect of oil on the stock market, there is an emerging literature on the financialization of commodities that could extend the short-term predictability of oil to other commodities. Indeed, the emergence of commodities as an asset class has increased the speculative trading in futures markets, especially for those commodities belonging to futures indices, such as the Goldman Sachs Commodity Index and the S&P Commodity Index. Tang and Xiong (2012) find that these indexed commodities have become increasingly correlated with oil since 2004 which makes us believe that these assets would also have some predictive power on stock returns.* Unfortunately, the sample period starting from 2004 is not long enough to do a formal study of predictability, so we leave the answer to this question for future research.

The rest of the paper is structured as follows. The next section reviews the relation between oil prices and both the business cycle and excess market returns, and also presents our variable. [Section 15.3](#) contains the results of in-sample predictability at a quarterly horizon. [Section 15.4](#) reports the out-of-sample tests. [Section 15.5](#) discusses the evidence of predictability at longer horizons. [Section 15.6](#) analyses the impact of predictability in the cross section of expected returns. [Section 15.7](#) concludes.

15.2 Oil Price, The Business Cycle and Stock Returns

This section presents the mechanisms that relate the oil prices to the business cycle and excess stock market returns. It also presents our variable.

15.2.1 Oil Price and the Macroeconomy

The study of the relationship between oil prices and the macroeconomy begins with the seminal work of Hamilton (1983). He uses Sims's (1980) six-variable VAR and bivariate VARs with quarterly data for the 1948–1980 period to show that oil price shocks strongly Granger-caused the U.S. GNP growth rate and the unemployment rate. He finds that an increase in the oil price is followed by four successive quarters of lower GNP growth rates. As shown by Ferderer (1996), the common transmission channels of oil price shocks to the real economy are: inflation, terms of trade (Huntington 2007), the capital utilization rate (Finn 2000) and increasing returns to scale in production (Aguilar-Conraria and Wen 2007).

Subsequent studies (Mork 1989, Lee *et al.* 1995, Hooker 1996, Hamilton 2008) note a weakening of this relationship when data from 1980 onwards are included (which coincides with OPEC's loss of control of the oil market).† This turned attention towards nonlinear relationships between the variables (Mork 1989, Ferderer 1996, Hamilton 1996, 2003). This evidence required new explanations to understand the asymmetric impact of oil on the economy. The most commonly cited asymmetric mechanisms are: monetary policy (Ferderer 1996, Bernanke *et al.* 1997, Balke *et al.* 2002, Hamilton and Herrera 2004, Leduc and Sill 2004), imperfect intersectoral mobility of factors (Lilien 1982, Hamilton 1988, Davis and

* See Masters (2008), Buyuksahin and Robe (2010), Etula (2010), Acharya *et al.* (2011), Basak and Pavlova (2012) and Singleton (2012) for more studies concerning the financialization of commodities and the effect of speculators on commodity prices.

† For example, Hamilton (2008) finds that while in the 1949–1980 period an increase of 10% in the oil price predicts that the GDP growth would be 2.9% slower four quarters later, this figure for the 1949–2005 period is only 0.7%.

Haltiwanger 2001, Lee and Ni 2002), investment irreversibility (Bernanke 1983), wage rigidities (Lee *et al.* 1995), and interest rates (Balke *et al.* 2002). Also, Carruth *et al.* (1998) document a strong and significant relationship between the U.S. rate of unemployment and oil prices.

Recently, Kliesen (2008) added to the standard regression the variable CFNAI (Chicago Fed National Activity Index), which is the first principal component of 85 monthly indicators of real economic activity, and finds that oil has a significant impact on the U.S. macroeconomic performance. In addition, Cologni and Manera (2009) find a negative influence of oil shocks on GDP growth, although they reject the hypothesis that real GDP growth has no effect on oil prices. Kilian (2008) finds empirical evidence that exogenous oil supply shocks have caused significant impacts on the GDP of G-7 countries.* Cologni and Manera (2008) observe that oil price shocks affect only the GDP in Italy and in the U.S., albeit temporarily. For almost all of the countries in their sample, oil shocks affect inflation and nominal exchange rates. Gronwald (2008) concludes that only oil shocks that exceed a certain threshold affect the real sector of the economy, while 'normal' positive shocks generate significant nominal impacts. On the contrary, Barsky and Kilian (2004) find inaccuracies in the explanations that support most theories and conclude that 'disturbances in the oil market are likely to matter less for U.S. macroeconomic performance than has commonly been thought'.

The above empirical evidence, although not free of debate, largely supports the existence of a significant relationship between oil price shocks and the business cycle in both economic and statistical terms.

15.2.2 Oil Price and the Financial Market

Contrary to what has occurred with the relationship between oil shocks and the macroeconomy, the linkage between oil and the financial markets has received little attention. Moreover, this reduced literature has not been conclusive about this relationship. For a deeper analysis, [Table 15.1](#) presents detailed information on the empirical studies reviewed in this section. The table also shows the multiple variables proxying for oil shocks that have been used in the literature.

At a country level, Jones and Kaul (1996) find that oil price shocks produce significant changes in stock returns. The relation can be explained by changes in cash flows and discount rates for the United States and Canada, however they find an overreaction in the United Kingdom and Japan. Driesprong *et al.* (2008) find that oil price shocks have significant predictive power in developed economies, but not for emerging countries. They find an initial underestimation by agents of the impacts of the shock that is slowly corrected later. Park and Ratti (2008) find evidence that oil price shocks have a significant negative impact on real returns of several net importing countries, unlike what occurs in Norway, a net exporter, where the impact is positive.

Following Kilian (2009), Kilian and Park (2009) break down oil shocks into three classes: supply, aggregate demand and specific demand (or precautionary demand) shocks. According to their results, these shocks explain 6%, 5% and 11% of the long-run variation of real stock returns, respectively. They do not find a significant response of stock returns to oil supply shocks, however they do find a positive response to global demand shocks and a negative response to precautionary demand shocks. Sector-level evidence suggests that the mechanism that transmits oil shocks to stock returns is through the demand for industrial products, and not, as widely believed, through the production costs of the firms. Apergis and Miller (2009) criticize the methodology used by Kilian (2009) and Kilian and Park (2009) due to the use of both stationary and non-stationary variables in their VAR specification. They differentiate the $I(1)$ variables and carry out the same breakdown as Kilian (2009), but using only $I(0)$ variables. Instead of including the equity returns in the same VAR, as in Kilian and Park (2009), Apergis and Miller (2009) use a second VAR with the three types of shocks and the stock market returns of each country. Using a sample composed of the G-7 Group and Australia, they conclude that the effects of oil shocks, although statistically significant, produce a minor impact on stock returns.

* This somehow contradicts Kilian (2009), who finds that the impacts of oil demand shocks are more significant than supply shocks.

TABLE 15.1 Literature on Financial Markets and Oil Prices. The Table Provides Information on Studies on Financial Markets and Oil Prices. The Variables Used are: Δf , the Log Returns on Oil Futures Prices; Δs , the Log Returns on Nominal Spot Prices; Δs_r , the Log Returns on Real Spot Prices; Sop , the Oil Prices Scaled by Volatility, Unexpected Changes in Oil Prices (Lee *et al.* 1995); $Nopi$, Net Oil Price Increases (Hamilton 1996, 2003); Vol , Rolling Volatility of Oil Price Changes; $x^+ = \max(0, x)$; $x^- = \min(0, x)$

Paper	Subject	Frequency	Stock Market	Data	
				Series	Variable
Huang <i>et al.</i> (1996)	Joint dynamics	Daily	S&P 500, industrial portfolios, oil companies	Heating and crude oil futures	Δf
Jones and Kaul (1996)	Market efficiency	Quarterly	Real country indexes	PPI oil and related products	Δs
Sadorsky (1999)	Joint dynamics	Monthly	S&P 500/CPI	PPI fuels/CPI	Δs_r , Δs_r^+ , Δs_r^- , Sop , Sop^+ , Sop^-
Ciner (2001)	Joint dynamics	Daily	S&P 500	Heating and crude oil futures	Δf
Driesprong <i>et al.</i> (2008)	Predictability	Monthly	Country indexes, sector indexes	Brent, WTI, Dubai, Arab Light, Brent futures, oil futures	Δs , Δf
Park and Ratti (2008)	Joint dynamics	Monthly	Real country indexes	Brent/PPI all commodities	Δs_r , Sop , $Nopi$, Δs_r^+ , Δs_r^- , Sop^+ , Sop^- , Vol
Kilian and Park (2009)	Oil demand and supply shocks	Monthly	Real CRSP value weighted, industry portfolios	EIA refiner acquisition cost/CPI	Δs_r
Apergis and Miller (2009)	Oil demand and supply shocks	Monthly	Real country indexes	EIA refiner acquisition cost/CPI	Δs_r

Some papers report a significant intratemporal correlation between oil shocks and portfolio returns. Nandha and Faff (2008) calculate a two-factor model that includes unexpected returns on the world market and the oil market. They find a significant negative impact of oil that is largely symmetrical on all sectors except for mining and oil/gas. On the other hand, Bachmeier (2008) regresses portfolio returns on shocks that are contemporaneous with the oil price. He finds that oil shocks have a significant negative impact on returns.

Finally, Huang *et al.* (1996) find that daily oil futures returns have no correlation with stock returns, except for returns on oil companies. Ciner (2001) tests for nonlinear Granger-causality of oil futures returns on stock market returns and finds a significant relationship. Using monthly data, Sadorsky (1999) finds that, in the U.S. market, the price of oil affects the financial market, but that the effect in the other direction is insignificant. He also finds that oil price shocks have an asymmetric effect on industrial production and real stock returns, with positive shocks having a greater impact than negative ones.

On the theoretical side, Wei (2003) builds a general equilibrium model to estimate the impact of an oil price shock on the value of a firm that faces investment irreversibility. His model predicts that an oil shock will have only a small impact. Consequently, he is unable to explain the massive decline in the stock market in 1974 after the oil shock of 1973.

15.2.3 Stock Returns and the Business Cycle

This section shows that stock market returns vary considerably with the business cycle. This suggests that a variable that anticipates the cycle, such as oil price shocks, may also do a good job at predicting stock returns.

Figure 15.1 shows the excess market returns from 1926Q3 to 2009Q4 and the shaded areas represent the NBER (National Bureau of Economic Research) recession periods.* The stock returns tend to be negative and grow during recessions, reaching peaks towards the end of each one. In fact, in our sample period the maximum return was reached by the end of the Great Depression of the 1930s. Nevertheless, this commonly cited contracyclical character of stock returns is evident only at the end of recessions; at the beginning of and during recessions, these returns are highly procyclical. For example, the minimum in the sample period was also reached during the Great Depression. In addition, expansions are characterized on average by positive returns, although they are less volatile than those generated during recessions.

NBER's business cycle dates enable the sample to be divided into four stages: expansion, peak, recession and trough. Table 15.2 shows the first two conditional sample moments for the excess stock returns. As can be seen in the table, the most frequent stage of the economy is expansion, and during this phase of the cycle the average return is positive (2.9%) and greater than its historical average (2.0%). On the contrary, during the recession stage the average excess return on the market has a similar magnitude as




FIGURE 15.1 Excess stock market returns, oil shocks and the business cycle, 1926Q3–2009Q4.

TABLE 15.2 Stock Market Returns and the Business Cycle. Conditional Excess Stock Returns for the Sample Period 1926Q3 to 2009Q4. Classification of States Based on NBER Business Cycle Dates

State	Frequency	Average Excess Return	Standard Deviation of Excess Returns
Expansion	252	2.9	9.1
Peak	15	−5.5	8.1
Recession	52	−3.0	17.3
Trough	15	12.3	10.0
Total	334	2.0	11.3

* For longer sample periods we use the stock returns and the risk-free rate available from Ken French's web page, which contains data from September 1926. Stock returns are from the value-weighted CRSP index and the risk-free rate is from the 3-month Treasury bills from Ibbotson Associates.

it does during the expansion stage, but with the opposite sign (-3.0%) and almost twice the volatility. Finally, during the peaks (troughs), expectations about the state of the economy are negative (positive) and therefore excess stock returns are highly negative (positive) in these stages.

15.2.4 Measuring Oil Price Shocks

Considering the evidence mentioned above, it seems natural to propose oil price shocks as a leading variable with the potential to forecast stock market returns. However, to maximize its predictive power, it is essential to consider only unanticipated changes in the oil price. Although oil spot returns have been a widely used variable in the literature (see Table 15.1), we do not consider them, because they contain some components that are clearly anticipated by market participants, such as the interest rate and the convenience yield.* One way to address this issue is to estimate unexpected oil changes from a model for the spot price dynamics, but this procedure still has disadvantages in that it depends on the model specification and that the information set used by the econometrician to estimate the conditional mean may not coincide with that of the market. Unexpected oil price changes can only be captured with an objective and precise estimate of the expected spot price in the future.

We measure unexpected oil price shocks by short-term futures returns on crude oil. Using Fama and French's (1987) methodology and cointegration tests, Switzer and El-Khoury (2007) show that oil futures prices have significant predictive power for future spot prices. Moreover, Ma (1989) and Kumar (1992) confirm that futures prices, in addition to being unbiased predictors of spot prices, exceed the predictive capacity of a random walk and a wide variety of models. This evidence suggests that unexpected changes in oil prices are correctly captured by our proposed variable.† Therefore, we assume that quarterly unexpected oil shocks are proxied by oil futures returns, i.e.

$$\Delta f(t) \equiv f^1(t) - f^4(t-3) \approx s(t) - \mathbb{E}_{t-3}[s(t)], \quad (15.1)$$

where $s(t)$ is the log oil spot price and $f^\tau(t)$ is the log oil futures price of a contract that matures in τ months.

Data on oil futures prices are from NYMEX, which began trading these contracts in March 1983. Therefore, our sample period goes from 1983Q2 to 2009Q4.‡ Figure 15.1 confirms that recessions are preceded by positive oil shocks, while during recessions these shocks are rapidly reversed. The great variability of the oil shocks is evidence of their predictive potential. Another visual characteristic is the low persistence of the series, which prevents it from being subject to the critiques of existing predictive variables. The optimal number of lags of our variable to be considered in this study was determined using the Akaike Information Criterion (AIC) in Ordinary Least Squares (OLS) regressions of the excess stock market return on lagged oil shocks. In this study we consider four lags of our variable, because this number minimizes the AIC (see Table 15.3). Interestingly, the number of lags coincides with that obtained in the macroeconomic literature (Hamilton 2003, 2008).

* The convenience yield of crude oil is a benefit for immediate ownership of physical units of the commodity attributed to the benefit of protecting regular production from temporary shortages of oil stocks (see, for example, Casassus and Collin-Dufresne (2005)).

† A direct way to test that futures prices are unbiased estimators of future spot prices is to verify the existence of a risk premium associated with this contract. In our sample, the t -stat of the null hypothesis that the average logarithmic futures return is zero is 1.13. Although this informal test only detects a constant risk premium, the results obtained by Switzer and El-Khoury (2007) suggest that, in oil futures, this finding is not incompatible with the existence of predictive power for spot prices.

‡ We use quarterly data to better capture the aggregate impact of oil shocks and make our results more compatible with the macroeconomic evidence. Moreover, this allows us to include in the analysis the consumption–wealth ratio, a variable which is only available at a quarterly frequency (Lettau and Ludvigson 2001a).

TABLE 15.3 Optimal Lags of Oil Price Shocks. OLS Regressions of Excess Stock Returns on Lags of Oil Price Shocks, 1983Q2–2009Q4. All Estimations Use the Full Sample and Include a Constant

Lags of oil price shocks	AIC
0	−4.884
1	−4.891
2	−4.867
3	−4.846
4	−4.893
5	−4.870
6	−4.844
7	−4.824
8	−4.803

15.2.5 Oil Price, The Business Cycle and Stock Returns

To show evidence of the relationship among the oil price, the business cycle and excess stock returns, we study the joint dynamics of these variables using a vector auto-regression analysis with four lags. Our proxy for the stock market is the value-weighted CRSP index, from which we obtain the quarterly returns on the market portfolio (R_m). We proxy the risk-free rate (R_f) with the 3-month constant-maturity treasury yields from the Federal Reserve Board of Governors. Following Cooper and Priestley (2009), we use the total Industrial Production Index (IP) from the Fed as a measure of output and a proxy for the business cycle. Table 15.4 shows the maximum likelihood estimates for the VAR(4) model. As is common in macroeconomic series, industrial production growth rate ($\Delta\%IP$) is the easiest series to predict (its R^2 is 51%) and its own lags have useful information for forecasting its future values. On the other hand, the predictive power of the excess stock return ($R_m - R_f$) on $\Delta\%IP$ is a clear signal that the financial market correctly anticipates future economic growth. Moreover, in our sample the Granger-causality of oil shocks (Δf) on economic growth is also verified, evidence that is in line with the macroeconomic studies mentioned above.

The table also shows that, for the excess stock returns, the lack of significance of its own lags suggests the stock market is efficient and that the initial underreaction and later correction documented by Driesprong *et al.* (2008) is not seen at quarterly frequency.* There is also evidence of inverse causality with the industrial production growth rate, which is probably the consequence of an adjustment process of previous expectations about actual economic growth. Furthermore, and as expected, oil shocks demonstrate a significant predictive power for excess stock returns, which will be explored in greater depth in the following sections.

Finally, the results reveal that oil shocks cannot be predicted with any of the lagged variables, which is evidence that our measure for oil shocks effectively captures unanticipated changes in this variable. This result implies that our variable is exogenous to the U.S. economy; however, it may not be so to global aggregate demand. Therefore, the critique of Kilian (2009) regarding the endogeneity of oil price shocks is still valid for our variable.

The following sections analyse and test the predictive power of oil price shocks for stock returns.

* If the adjustment process of returns after an oil shock is slow, this should be evidenced to some degree in the excess stock return, either positively if the adjustment is gradual or negatively if it is excessive and requires future reversals.

TABLE 15.4 VAR Estimation Results, 1983Q2–2009Q4. Maximum Likelihood Estimates of the VAR(4) Model for the Rate of Growth of Industrial Production ($\Delta\%IP(t)$), Excess Stock Market Returns ($R_m(t) - R_f(t)$) and Log Returns on Crude Oil Futures ($\Delta f(t)$). Asymptotic t -Stat in Parentheses. Granger Causality was Tested Using the Asymptotic Wald Test. The *All* Row at the Bottom of the Table Refers to All Coefficients Except the Constant

	$\Delta\%IP(t)$	$R_m(t) - R_f(t)$	$\Delta f(t)$
Constant	0.000 (0.45)	0.011 (1.20)	0.020 (0.88)
$\Delta\%IP(t - 1)$	0.276 (3.04)	0.956 (1.25)	-0.205 (-0.11)
$\Delta\%IP(t - 2)$	-0.027 (-0.26)	-0.852 (-0.99)	0.404 (0.19)
$\Delta\%IP(t - 3)$	-0.113 (-1.10)	-1.215 (-1.40)	0.460 (0.22)
$\Delta\%IP(t - 4)$	0.347 (3.70)	2.327 (2.93)	-1.194 (-0.62)
$R_m(t - 1) - R_f(t - 1)$	0.061 (5.59)	0.038 (0.41)	0.269 (1.19)
$R_m(t - 2) - R_f(t - 2)$	0.030 (2.52)	-0.069 (-0.68)	0.191 (0.77)
$R_m(t - 3) - R_f(t - 3)$	0.005 (0.42)	-0.038 (-0.37)	0.268 (1.07)
$R_m(t - 4) - R_f(t - 4)$	0.026 (2.12)	0.019 (0.19)	-0.087 (-0.35)
$\Delta f(t - 1)$	0.010 (2.09)	0.065 (1.59)	0.035 (0.35)
$\Delta f(t - 2)$	-0.006 (-1.22)	-0.044 (-1.04)	-0.183 (-1.79)
$\Delta f(t - 3)$	-0.012 (-2.35)	-0.004 (-0.10)	0.027 (0.26)
$\Delta f(t - 4)$	-0.005 (-1.00)	-0.119 (-2.74)	-0.057 (-0.54)
R^2	0.51	0.18	0.06
$\overline{R}^2$	0.44	0.07	-0.06
p -Value Granger causality test			
<i>All</i>	0.00	0.03	0.85
$\Delta\%IP$	0.00	0.04	0.98
$R_m - R_f$	0.00	0.94	0.53
Δf	0.01	0.04	0.49

15.3 Short-Horizon Predictability of Stock Returns

This section focuses on the in-sample predictability of stock returns at a quarterly horizon. The predictive performance of our variable is evaluated and compared with the performance of the following variables: the risk-free rate (Campbell 1987), the log dividend–price ratio (Fama and French 1988), the consumption–wealth ratio (Lettau and Ludvigson 2001a) and the output gap (Cooper and Priestley 2009). The log dividend–price ratio ($d - p$) was calculated from the value-weighted CRSP index using

the methodology described by Ang and Bekaert (2007). The consumption–wealth ratio (*cay*) and its individual components are from Martin Lettau’s web page and sampled at a quarterly frequency. The output gap (*gap*) is constructed using the total Industrial Production Index.* The output *gap* is estimated with data from 1948Q1 to replicate the series from Cooper and Priestley (2009). The variables *cay*(*t*) and *gap*(*t*) are assumed to be known at the start of time *t* + 1, and therefore can be used to forecast excess stock returns. We omit any complication due to the look-ahead bias in these variables and the normal delay in the publication and subsequent revisions of these and other macroeconomic series.

Table 15.5 presents the main statistics of the predictive variables and Figure 15.2 provides graphical evidence. The upper panel of Table 15.5 shows that our variable (Δf) and $d - p$ exhibit more volatility than the variables R_f , *cay* and *gap*. Unlike the other predictors, Δf shows very low persistence; in fact, its first-order serial correlation is as low as that of $R_m - R_f$. The lower panel of the table shows that the intra-temporal correlation of our variable with the excess stock return is very low, rejecting Δf as a possible




FIGURE 15.2 Excess stock market returns and predictor variables, 1983Q2–2009Q4.

TABLE 15.5 Statistics for 1983Q2–2009Q4. Autocorrelation is the First-Order Serial Correlation

	$R_m - R_f$	Δf	R_f	$d - p$	<i>cay</i>	<i>gap</i>
Average	0.016	0.021	0.012	−3.781	0.005	−0.014
Standard deviation	0.087	0.195	0.006	0.386	0.019	0.063
Autocorrelation	0.031	0.031	0.972	0.974	0.910	0.977
Correlation matrix						
	$R_m - R_f$	Δf	R_f	$d - p$	<i>cay</i>	<i>gap</i>
$R_m - R_f$	1.000					
Δf	−0.039	1.000				
R_f	0.006	−0.035	1.000			
$d - p$	0.139	−0.101	0.524	1.000		
<i>cay</i>	−0.108	−0.100	0.348	0.472	1.000	
<i>gap</i>	−0.121	0.111	−0.210	−0.847	−0.601	1.000

* Cooper and Priestley (2009) define four methods for calculating the output gap, but the main one is the quadratic version of *gap* (based on its greater correlation with procyclical variables), which we also use here. The output gap is calculated from the following regression: $ip(t) = a + bt + ct^2 + gap(t)$, where $ip(t)$ is the log industrial production, t is a time trend and $gap(t)$ is the error term.

pricing factor. Both the low persistence of our variable and its low correlation with excess stock returns discard the idea that the forecastability of stock returns from oil returns might be due to a common factor driving both variables. Another important characteristic of our variable is its low correlation with other predictive variables, which suggests that Δf contains business cycle information not captured by the existing predictors. On the contrary, the existing variables show high levels of correlation (in absolute terms) among themselves, revealing the presence of redundant information.

Next we turn our attention to evaluating the predictive performance of our variable and the other variables considered here. We begin with the evidence of in-sample predictability at a quarterly horizon, for which we estimate the following regression:

$$R_m(t) - R_f(t) = \alpha + \beta'X(t-1) + \varepsilon(t), \quad (15.2)$$

where $X(t-1)$ is a vector of known predictors at $t-1$ and β its associated coefficient vector. It should be emphasized that $X(t-1)$ can include one variable, several variables or several lags of the same variable.

Table 15.6 shows the OLS results of Equation (15.2). We report asymptotic t -stats and Wald tests that correct for serial correlation and heteroscedasticity using Newey and West (1987).^{*} The columns present the results for each predictor variable. The results indicate that our variable has the best in-sample predictive performance with an $\bar{R}^2$ of 6%. The Wald test corresponds to the null hypothesis that all the

TABLE 15.6 Predictive Regressions for Excess Stock Returns, 1983Q2–2009Q4. OLS Regressions for Excess Stock Returns on the Predictor Variables in the First Row. All Tests are Based on Covariance Matrices of Coefficients Corrected For Heteroscedasticity and Serial Correlation Using Newey and West (1987). Lag Length in the Newey–West Estimator is $\text{floor}[4 \cdot (T/100)^{2/9}]$, Where $\text{Floor}[X]$ Denotes the Integer Part of X (Newey and West 1994). Asymptotic t -Stat in Parentheses. At the Bottom of the Table, the P -Value is for the Asymptotic Wald Test and *All* Refers to All Coefficients Except the Constant

	Δf	R_f	$d - p$	<i>cay</i>	<i>gap</i>
Constant	0.019 (2.59)	0.013 (0.55)	0.154 (2.04)	0.014 (1.52)	0.015 (1.66)
$\Delta f(t-1)$	0.068 (1.95)				
$\Delta f(t-2)$	-0.035 (-0.53)				
$\Delta f(t-3)$	-0.024 (-0.48)				
$\Delta f(t-4)$	-0.115 (-2.68)				
$R_f(t)$		0.403 (0.25)			
$d(t-1) - p(t-1)$			0.036 (1.75)		
<i>cay</i> ($t-1$)				0.729 (2.09)	
<i>gap</i> ($t-1$)					-0.219 (-1.58)
R^2	0.10	0.00	0.02	0.02	0.02
$\bar{R}^2$	0.06	-0.01	0.01	0.01	0.02
p -Value Wald <i>All</i>	0.02	0.80	0.08	0.04	0.11

^{*} Following Newey and West (1994), we chose lag length equal to $\text{floor}[4 \cdot (T/100)^{2/9}]$, where $\text{floor}[x]$ denotes the integer part of x .

coefficients in Equation (15.2) are zero, except the constant. This statistic is highly significant for the oil price shocks (p -value of 2%). The annual cumulative impact of our variable is -0.106 and is calculated from the sum of the coefficients corresponding to the four lags of Δf . This impact is also economically significant. To see this, consider a one-time increase in Δf of one standard deviation (19.5% in our sample, see Table 15.5). This change leads to a 2.1% decrease in expected quarterly excess returns on the value-weighted CRSP index ($-0.106 \times 19.5\% = -2.1\%$), which is equivalent to 29.6% ($2.1\%/7.0\% = 29.6\%$) of the historical average annual excess return.

The dynamics of the distributed lags can be explained with aggregate demand shocks (Kilian and Park 2009). A positive shock to the global demand for industrial commodities produces both a direct positive impact and an indirect negative one in the financial market. The direct impact manifests as an increase in both the oil price and economic growth with a consequent positive stock return. Increased economic growth pushes the oil price even higher, and thus indirectly affects negatively the future expected economic growth and expected stock returns. The final impact will depend on the relative magnitudes of both impacts.* The direct impact is initially stronger, which explains the positive significant effect of the first lag of Δf . Later, the indirect negative impact of the second and third lags begins to gather some strength, although not of a sufficient magnitude to cancel out the initial positive impact (see the sign and low significance of the following two lags). One year after the unexpected aggregate demand shock, the indirect effect becomes dominant; in other words, the high price of oil causes a deceleration in the economy. This is manifested by the negative and significant coefficient of $\Delta f(t-4)$ which is also responsible for the cumulative negative impact reported.

The third column of Table 15.6 shows the results for the interest rate as a predictor. This variable has the worst predictive performance in our sample, with a $\bar{R}^2$ of -0.01 . Contrary to the findings of previous studies (Campbell 1987), the coefficient that accompanies this variable is positive, although not significant. This poor performance is associated with the low volatility of this variable in our sample. Its standard deviation is 0.006 (see Table 15.5), a value well below the 0.032 reported by Ang and Bekaert (2007) for 1935Q2–2001Q4. The last three columns contain the results for the $d-p$, cay and gap variables. All of these have intuitive signs, although their coefficients are lower in absolute value than those reported in previous studies (see Ang and Bekaert (2007), Lettau and Ludvigson (2001a) and Cooper and Priestley (2009), respectively).† Moreover, the R^2 statistic for these variables (all about 2%) suggest that they have similarly poor predictive power. Among these variables, cay has a better performance, being the only one significant at the 5% level.

The previous results suggest that the oil shocks have a significant in-sample forecasting power for short horizons. We now see whether our variable is robust to the inclusion of the other predictors considered here. To evaluate this, we estimate the following extended predictive regression:

$$R_m(t) - R_f(t) = \alpha + \sum_{j=1}^4 \beta_j \Delta f(t-j) + \theta' Z(t-1) + \epsilon(t), \quad (15.3)$$

where $Z(t-1)$ is a vector of predictor variables and θ its associated coefficients vector. Lack of robustness in the predictive power of our variable should be reflected in changes in sign and/or loss of significance in the coefficients that accompany the lags of oil shocks.

The results of the estimation of Equation (15.3) are presented in Table 15.7. The last row contains the p -value of an asymptotic Wald test for the combined null hypothesis that all the coefficients associated with the lags of our variable are zero. The columns show the results of including each of the other variables in the predictive regression of Δf , while the inclusion of all of them is considered in the last

* Note from our VAR analysis in the previous section that the signs of the four lags of oil shocks are the same for both industrial production and excess stock returns (see Table 15.4). This result gives support to our aggregate demand based explanation.

† According to Stambaugh (1999), the estimation of the β coefficient in a predictive regression of $d-p$ is biased. This bias further reinforces our conclusion regarding this variable's lack of forecasting power.

TABLE 15.7 Predictive Regressions: Additional Controls, 1983Q2–2009Q4. OLS Regressions for Excess Stock Returns on the Predictor Variables in the First Row. All Tests are Based on Covariance Matrices of Coefficients Corrected for Heteroscedasticity and Serial Correlation Using Newey and West (1987). Lag Length in the Newey–West Estimator is $\text{Floor}[4. (T/100)^{2/9}]$, Where $\text{floor}[x]$ Denotes The Integer Part of x (Newey and West 1994). Asymptotic t -Stat in Parentheses. At the Bottom of the Table, the p -Value is for the Asymptotic Wald Test, *All* Refers to All Coefficients Except the Constant and Δf Refers to the Coefficients Associated With the Four Lags of That Variable

	$\Delta f \& R_f$	$\Delta f \& d - p$	$\Delta f \& cay$	$\Delta f \& gap$	$\Delta f \& All$
Constant	0.012 (0.57)	0.140 (2.02)	0.015 (1.81)	0.017 (2.01)	0.269 (1.03)
$\Delta f(t-1)$	0.067 (1.96)	0.075 (2.13)	0.075 (2.07)	0.076 (2.20)	0.078 (1.87)
$\Delta f(t-2)$	-0.036 (-0.53)	-0.026 (-0.41)	-0.029 (-0.47)	-0.025 (-0.38)	-0.028 (-0.47)
$\Delta f(t-3)$	-0.026 (-0.49)	-0.017 (-0.33)	-0.023 (-0.45)	-0.013 (-0.26)	-0.026 (-0.53)
$\Delta f(t-4)$	-0.116 (-2.72)	-0.108 (-2.50)	-0.114 (-2.67)	-0.105 (-2.33)	-0.117 (-2.54)
$R_f(t)$	0.650 (0.45)				-1.708 (-0.64)
$d(t-1) - p(t-1)$		0.032 (1.68)			0.061 (1.01)
$cay(t-1)$			0.784 (2.00)		0.985 (1.78)
$gap(t-1)$				-0.172 (-1.50)	0.303 (0.93)
R^2	0.10	0.11	0.12	0.11	0.14
$\bar{R}^2$	0.05	0.07	0.08	0.06	0.06
p -Value Wald <i>All</i>	0.02	0.00	0.00	0.00	0.00
p -Value Wald Δf	0.01	0.01	0.01	0.02	0.02

column. First, given the low correlation of our variable with the others, the forecasting power of our variable remains intact. According to the Wald test for the coefficients of Δf , in all of the estimates these coefficients keep their joint significance. In addition, their signs and individual significance remain the same, and they are roughly the same size. Second, consistent with previous results, the greatest increase in predictive power is reached when our variable is used in combination with *cay*, obtaining an $\bar{R}^2$ of 8%. Third, contrary to what occurs with the oil shock coefficients, the results of the last column provide evidence of great instability in the predictive power of the other variables. None of these are significant. The coefficients of R_f and *gap* experience changes in sign and those of $d - p$ and *cay* vary dramatically in size. Of course, this evidence is consistent with the high correlation between these variables reported in Table 15.5.

In summary, this section demonstrates that our variable has significant and robust short-horizon in-sample forecasting power for stock returns. Out-of-sample predictability, also at a quarterly horizon, is considered in the next section.

15.4 Out-of-Sample Predictability of Stock Returns

In-sample predictive performance is essential for establishing the existence of predictability. However, in order for a predictor to be used by an investor, it must also demonstrate a good out-of-sample performance. That is, a predictive variable must be able to forecast excess returns reasonably well with information available at the time of the forecast, which is not guaranteed by the in-sample tests of Equations (15.2) and (15.3), as the coefficients are estimated using the full sample.

Welch and Goyal (2008) conclude that it is very difficult to find variables with short-horizon out-of-sample forecasting power that outperform the average excess return in a recent sample period. In fact, when considering the sub-period from 1965 to 2005, they only find variables that outperform the prevailing historical average return at a five-year horizon. Although the out-of-sample predictive performance can be increased by imposing certain restrictions, as shown by Campbell and Thompson (2008), here we choose to keep the simplicity and linearity of the predictive model. We also test the other variables' out-of-sample performance despite their poor in-sample predictive power.

In order to contribute to this discussion, we compare forecasts from nested linear models to determine whether each variable has predictive content for stock returns. The prevailing historical average of excess stock returns is used as a benchmark. Therefore, we define the following benchmark and competing models:

$$\text{benchmark: } R_m(t+1) - R_f(t+1) = \alpha_1 + u_1(t+1), \quad (15.4)$$

$$\text{competing: } R_m(t+1) - R_f(t+1) = \alpha_2 + \beta'X(t) + u_2(t+1), \quad (15.5)$$

where the coefficients α_1 , α_2 and β are estimated recursively. The sample size T is divided into in-sample and out-of-sample portions. Q is defined as the minimum number of observations used for estimating the coefficients and $P \equiv T - Q$ denotes the maximum number of one-step-ahead predictions. Thus, forecasts of $R_m(t+1) - R_f(t+1)$, $t = Q, \dots, T-1$, are generated recursively using the two linear models in Equations (15.4) and (15.5), where all coefficients are re-estimated with new observations as forecasting moves forward through time.

Our assessment of out-of-sample predictability involves three metrics that are presented with more detail in Appendix A. The first approach is the forecast encompassing test of Clark and McCracken (2001) that verifies whether the competing forecasts incorporate any useful information absent in the benchmark forecasts. Their statistic ENC-NEW tests the null hypothesis that the benchmark model encompasses the competing model. This test suffers from a look-ahead bias, since ENC-NEW depends on a parameter that is estimated with the full sample. The second metric we use is the MSE-F test developed by McCracken (2007) that compares the predictive accuracy between nested models measured as the mean squared forecasting errors (MSE). Although this test does not have the best small-sample properties, it enables testing of out-of-sample predictive power under the same conditions that an investor faces in reality. The final measure of out-of-sample forecasting performance is the out-of-sample R^2 , R_{Os}^2 , proposed by Campbell and Thompson (2008). Following Welch and Goyal (2008), we use bootstrap in order to address size distortions in the t -stat for long horizons (Ang and Bekaert 2007) and to make the correct inference from nested out-of-sample predictability tests (Clark and McCracken 2005). We also consider covariance matrices of coefficients corrected for heteroscedasticity and autocorrelation that arise from the use of distributed lags and overlap of returns (Newey and West 1987).

In general, when performing out-of-sample tests in a small sample, there is a trade-off with the number of in-sample observations that is hard to resolve. On the one hand, the objective is to use a relatively large in-sample proportion of the sample (Q/T), so that the out-of-sample forecasts are done with estimates which are as similar as possible to those obtained with the full sample. But at the same time, as suggested by the results of Inoue and Kilian (2005), the out-of-sample proportion (P/T) must be large enough to prevent significant differences in power between the in-sample and out-of-sample tests. Thus, to achieve a reasonable level of power without producing excessive forecasting errors at the beginning of the out-of-sample sub-period, the optimal choice should be around $\pi = P/Q = 1$. However, to make our test more rigorous, we choose 1997Q4, which is when the Asian crisis hit the U.S. economy, as the starting point for the out-of-sample sub-period. That is, given that our adjusted (for lags) sample spans the period 1984Q2–2009Q4, our choice implies the following sample portions: $Q = 54$, $P = 49$ and $\pi = 0.91$.

The results of the out-of-sample tests are presented in Table 15.8. All of the tests coincide in that our variable is the only one with out-of-sample forecasting power for the excess stock returns at a 10% significance level.*

Although the *cay* variable is marginally significant according to the ENC-NEW test (its bootstrapped *p*-value is 10.3%), as explained above, the only test that measures forecasting power under the effective conditions faced by a potential investor is the MSE-F test. Our variable (Δf) has the highest and most significant R^2_{OS} among all of the variables considered; however, the size of this statistic is only 1.2%. The table also shows that, given the wide differences between the bootstrapped and asymptotic critical values, controlling for considerations of small sample and differences in the relative out-of-sample portion (i.e. π) is essential for a reliable inference, especially when highly persistent predictor variables are used. Also, as a consequence of the close relationship between the MSE-F and the R^2_{OS} statistics shown in Appendix A, the inference using the bootstrap method produces the same results for both tests.

Finally, the low R^2_{OS} for all the predictors is evidence that out-of-sample forecasting of stock returns has become an increasingly difficult challenge in recent times (one of the main points emphasized by Welch and Goyal (2008)). The forecasting ability of a predictor variable depends exclusively on its capacity to successfully summarize the conditioning information used by the market participants, which has become increasingly complex.

TABLE 15.8 Out-of-Sample Predictability Tests, 1983Q2–2009Q4. Out-of-Sample Tests of Stock Return Predictability. Each Column Reports the Results Using the Predictor from the First Row. Out-of-Sample Period is From 1997Q4 to 2009Q4. ENC-NEW, MSE-F and R^2_{OS} Statistics are Described In Equations (A.6), (A.7) and (A.8). Asymptotic Critical Values for the ENC-NEW Test are From Table 1 of Clark and Mccracken (2001) Using $\pi = 1.0$. Asymptotic Critical Values for the MSE-F Test are From Table 4 of Mccracken (2007) Using $\pi = 1.0$. Bootstrapped *p*-Values and Critical Values are Based on the Methodology of Clark and Mccracken (2005)

	Δf	R_f	$d - p$	<i>cay</i>	<i>gap</i>
ENC-NEW					
Sample value	2.218	−0.555	0.202	1.630	0.063
0.10 Asymptotic critical value	2.169	0.984	0.984	0.984	0.984
0.05 Asymptotic critical value	3.007	1.584	1.584	1.584	1.584
Bootstrapped <i>p</i> -value	0.077	0.791	0.278	0.103	0.505
0.10 Bootstrapped critical value	1.939	1.325	1.325	1.653	2.163
0.05 Bootstrapped critical value	2.970	2.033	2.276	2.559	3.294
MSE-F					
Sample value	0.603	−1.287	−0.034	0.261	−0.750
0.10 Asymptotic critical value	0.545	0.751	0.751	0.751	0.751
0.05 Asymptotic critical value	1.809	1.548	1.548	1.548	1.548
Bootstrapped <i>p</i> -value	0.081	0.655	0.172	0.197	0.463
0.10 Bootstrapped critical value	0.443	0.789	0.431	1.335	1.816
0.05 Bootstrapped critical value	1.420	1.860	1.268	2.237	3.006
R^2_{OS}					
Sample value	0.012	−0.027	−0.001	0.005	−0.016
Bootstrapped <i>p</i> -value	0.081	0.655	0.172	0.197	0.463
0.10 Bootstrapped critical value	0.009	0.016	0.009	0.027	0.036
0.05 Bootstrapped critical value	0.028	0.037	0.025	0.044	0.058

* Our results could be distorted by a possible look-ahead bias in the *cay* and *gap* variables, since for the sake of simplicity, both were calculated using the full sample of observations. Strictly speaking, in an out-of-sample test, forecasting of the excess stock returns implies the estimation of the coefficients with data only up to the prevailing quarter.

15.5 Long-Horizon Predictability of Stock Returns

This section examines the in-sample predictability of stock returns at longer horizons. The evidence presented below is based on the following long-horizon regression:

$$R_m^h(t+h) - R_f^h(t+h) = \alpha_h + \beta_h' X(t) + \varepsilon^h(t+h), \quad (15.6)$$

where $R_i^h(t+h) = \prod_{j=1}^h (1 + R_i^1(t+j)) - 1$ is the h -period return for asset $i = m, f$ and $R_i^1(t+j)$ is the respective one-period return from time $t+j-1$ to $t+j$.

The evidence of long-run predictability has been the subject of a great deal of criticism. According to Boudoukh *et al.* (2008) long-horizon regressions are misleading for persistent predictors.* That is, long-horizon regressions associated with the variables R_f , $d - p$, cay and gap cannot show anything different to what was shown in Section 15.3 due to their high persistence (see Table 15.5). On the other hand, given the almost null persistence of our variable, the Δf shocks are absolutely short-lived. Therefore, at the most our variable could have forecasting power for stock returns over a one-year horizon.

Following Kilian (1999), we adapt the bootstrap algorithm described in Appendix A to support the inference from the long-horizon regressions presented here. In addition, to evaluate the impact of the findings of Boudoukh *et al.* (2008) regarding persistent predictors, we report both expected regression coefficient and the R^2 statistic at the h th horizon conditional on their one-period counterpart, under the null of non-predictability. The expected coefficients are given by the following equations:

$$\mathbb{E}[\hat{\beta}_h | \hat{\beta}_1 = \hat{\beta}_1^*] = \left(1 + \frac{\rho(1-\rho^{h-1})}{1-\rho} \right) \hat{\beta}_1^*, \quad (15.7)$$

$$\mathbb{E}[R_h^2 | R_1^2 = R_1^{2*}] = \frac{(1 + \{\rho(1-\rho^{h-1})\} / (1-\rho))^2}{h} R_1^{2*}, \quad (15.8)$$

where $\hat{\beta}_1^*$ is the actual estimate of the regression coefficient in Equation (15.6) for $h = 1$, ρ is the first-order serial correlation coefficient of the predictor variable and R_1^{2*} is the actual estimate of the R^2 statistic for $h = 1$, also from Equation (15.6).

Our analysis covers return horizons up to five years ($h = 20$). For each horizon we consider the same number of observations and the same sample period 1984Q1–2004Q4 (i.e. 84 observations for each horizon). Table 15.9 presents the OLS estimates of Equation (15.6) for $h = 3, 4, 8, 12, 16, 20$. We report bootstrapped p -values for t -stats and Wald tests that correct for serial correlation and heteroscedasticity using Newey and West (1987). Because of overlapping observations, we increase the lag length for the Newey–West estimator by $h - 1$. The table shows that our variable has significant forecasting power for stock returns up to an horizon of three quarters. For $h = 3$ the signs of the coefficients associated with Δf lags are the same as those for $h = 1$, the $\bar{R}^2$ is 7% and the bootstrapped p -value for the test that all coefficients for oil lags are zero is 0.02. At longer horizons, as expected, there is no evidence of predictability with our variable.

Regarding the other variables, Table 15.9 shows that, at the 10% significance level, none of the variables demonstrates forecasting power for stock returns at long horizons. The cay variable is marginally significant at a five-year horizon; however, given the absence of predictive ability at other horizons, it is very likely that this result can be explained by the look-ahead bias. The variable R_f has so little forecasting power that despite its high persistence, it does not exhibit the pattern predicted

* Other problems documented for long-horizon regressions are: (1) serial correlation in residuals induced by the overlap of observations; (2) inefficient use of the data that provide spurious forecasts about the dynamics of variables, especially for non-exogenous predictors (Campbell 1991); and (3) by aggregating returns, long-horizon regressions invalidate the inference from standard asymptotic methods (Valkanov 2003).

TABLE 15.9 Long-Horizon Predictability, 1983Q2–2009Q4. OLS Regressions of $R_m^h(t+h) - R_f^h(t+h)$ on the Predictor Variables Using the Same Number of Observations. All Tests are Based on Covariance Matrices of Coefficients Corrected for Heteroscedasticity and Serial Correlation (Newey and West 1987). Lag Length in the Newey–West Estimator is $\text{Floor}[4 \cdot (T/100)^{2/9}] + (h-1)$, Where $\text{Floor}[x]$ Denotes the Integer Part of x . Bootstrapped p -Values for t -Stats in Parentheses. $\mathbb{E}[\hat{\beta}_h | \hat{\beta}_1 = \hat{\beta}_1^*]$ and $\mathbb{E}[R_h^2 | R_1^2 = R_1^{2*}]$ are Described in Equations (15.7) and (15.8), Respectively. p -Value Wald All is the Bootstrapped p -Value for the Asymptotic Wald Test that all Coefficients Except the Constant are Zero

Forecast horizon in quarters (h)	$h = 3$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
$\Delta f(t)$	0.153 (0.01)	0.019 (0.44)	-0.019 (0.45)	0.105 (0.33)	0.112 (0.35)	-0.143 (0.28)
$\Delta f(t-1)$	-0.051 (0.27)	-0.057 (0.31)	-0.057 (0.39)	0.029 (0.47)	-0.147 (0.30)	-0.323 (0.18)
$\Delta f(t-2)$	-0.081 (0.24)	-0.142 (0.18)	-0.128 (0.28)	-0.105 (0.35)	-0.051 (0.42)	-0.380 (0.14)
$\Delta f(t-3)$	-0.208 (0.05)	-0.146 (0.14)	-0.154 (0.21)	-0.119 (0.31)	-0.138 (0.30)	-0.432 (0.08)
R^2	0.12	0.05	0.02	0.01	0.01	0.04
$\bar{R}^2$	0.07	0.00	-0.03	-0.04	-0.04	0.00
p -Value Wald All	0.02	0.58	0.82	0.75	0.71	0.45
$R_f(t+1)$	0.682 (0.39)	0.926 (0.38)	1.129 (0.38)	-0.084 (0.57)	-3.804 (0.47)	5.304 (0.34)
$\mathbb{E}[\hat{\beta}_h \hat{\beta}_1 = \hat{\beta}_1^*]$	-0.125	-0.165	-0.317	-0.456	-0.585	-0.702
R^2	0.00	0.00	0.00	0.00	0.00	0.00
$\mathbb{E}[R_h^2 R_1^2 = R_1^{2*}]$	0.00	0.00	0.00	0.00	0.00	0.00
$\bar{R}^2$	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
p -Value Wald All	0.87	0.87	0.90	0.99	0.78	0.85
$d(t) - p(t)$	0.083 (0.13)	0.111 (0.14)	0.220 (0.18)	0.324 (0.18)	0.440 (0.18)	0.672 (0.14)
$\mathbb{E}[\hat{\beta}_h \hat{\beta}_1 = \hat{\beta}_1^*]$	0.078	0.102	0.190	0.266	0.331	0.388
R^2	0.06	0.08	0.13	0.14	0.17	0.25
$\mathbb{E}[R_h^2 R_1^2 = R_1^{2*}]$	0.05	0.06	0.10	0.14	0.16	0.17
$\bar{R}^2$	0.04	0.06	0.12	0.13	0.16	0.24
p -Value Wald All	0.28	0.29	0.36	0.36	0.36	0.28
$cay(t)$	2.501 (0.26)	3.436 (0.28)	7.493 (0.30)	12.588 (0.36)	18.904 (0.25)	25.065 (0.10)
$\mathbb{E}[\hat{\beta}_h \hat{\beta}_1 = \hat{\beta}_1^*]$	2.867	3.662	6.212	7.986	9.222	10.083
R^2	0.07	0.10	0.21	0.30	0.44	0.49
$\mathbb{E}[R_h^2 R_1^2 = R_1^{2*}]$	0.09	0.11	0.16	0.17	0.17	0.17
$\bar{R}^2$	0.06	0.09	0.20	0.30	0.43	0.48
p -Value Wald All	0.27	0.30	0.32	0.39	0.26	0.11
$gap(t)$	-0.444 (0.60)	-0.578 (0.62)	-1.548 (0.54)	-2.635 (0.45)	-4.161 (0.33)	-5.938 (0.25)
$\mathbb{E}[\hat{\beta}_h \hat{\beta}_1 = \hat{\beta}_1^*]$	-0.345	-0.454	-0.862	-1.229	-1.558	-1.853
R^2	0.04	0.05	0.14	0.21	0.34	0.43
$\mathbb{E}[R_h^2 R_1^2 = R_1^{2*}]$	0.02	0.03	0.05	0.07	0.08	0.09
$\bar{R}^2$	0.02	0.03	0.13	0.20	0.33	0.43
p -Value Wald All	0.64	0.66	0.57	0.47	0.34	0.26

by Boudoukh *et al.* (2008). Instead, its regression coefficient changes sign through horizons and its $\bar{R}^2$ is always negative. Finally, the other persistent predictor variables ($d - p$, *cay* and *gap*) effectively follow the pattern predicted by Boudoukh *et al.* (2008). That is, both regression coefficients and R^2 are always (in absolute value) growing with the horizon. In fact, both the regression coefficients and the R^2 values look pretty much like those predicted by Equations (15.7) and (15.8). These results reinforce the scepticism with which the evidence of predictability with highly persistent variables should always be viewed.

15.6 Implications for the Cross Section of Expected Returns

In this section we study the cross sectional implications of our variable. Despite the clear relationship between predictability, time-varying risk premium and conditional asset pricing models, there is an important methodological asymmetry between studies on predictability and those on conditional asset pricing. Papers on predictability have only sought to prove that the equity risk premium is variable and do not consider cross sectional tests. On the other hand, studies on conditional asset pricing only present evidence of in-sample predictability and generally lack the empirical rigor of the predictability studies.

In our opinion, if a variable exhibits significant forecasting power for stock market returns, it should also be tested as a conditioning variable to explain the cross section of expected stock returns. For this reason and for greater robustness, we offer empirical evidence in the cross-section of expected returns through CAPM and CCAPM models conditioned on our variable.* In addition, the empirical performance of our conditional model is compared with the performance of the same models conditioned on the other predictor variables considered in the previous sections, as well as the following unconditional models: CAPM (Sharpe 1964), CCAPM (Breedon and Litzenberger 1978) and the three-factor FF model (Fama and French 1993).

The standard equilibrium condition for any asset pricing model is given by

$$1 = \mathbb{E}[M(t+1) \cdot (1 + R_i(t+1)) | \Omega(t)], \quad (15.9)$$

where $M(t+1)$ is the stochastic discount factor or pricing kernel and $\Omega(t)$ is the agent's information set at time t . We assume that the pricing kernel is exponentially affine on the pricing factor:

$$M(t+1) = \exp(b^0(t) + b^1(t)F(t+1)), \quad (15.10)$$

where $b^i(t) = b_0^i + b_1^i z_1(t) + \dots + b_L^i z_L(t)$, $i = 0, 1$, and $F(t+1)$ is the risk factor. We use the aggregate market return, $R_m(t+1)$, as the risk factor for the CAPM models, and the real per-capita consumption growth rate, $\Delta c(t+1)$, for the CCAPM models. We allow for more than one conditioning variable, in particular we use $L = 4$ for our models, whereas $L = 1$ for conditional models on other variables. The exponentially affine form in Equation (15.10), as opposed to the standard linear assumption, is to keep the pricing kernel positive and to avoid the large negative values implied in linear models (Nagel and Singleton 2011).

* Multiple conditioning variables have been used in the literature. For example, Lettau and Ludvigson (2001b) propose CAPM and CCAPM models conditioned on the *cay* variable. Lustig and Van Nieuwerburgh (2005) derive a conditional CCAPM on the housing collateral ratio (ratio of housing wealth to human wealth). Santos and Veronesi (2006) present a conditional CAPM on the labour income to consumption ratio (fraction of consumption funded by labour income). In these studies the evidence of predictability is only an intermediate step that enables the authors to fulfil the first requisite needed for the existence of a conditional asset pricing model, or in other words, that the equity risk premium is variable. Then, given that the predictability condition is satisfied, there should be a conditional model that correctly values the cross section of returns using the proposed predictor.

As noted by Lewellen *et al.* (2010), conditional asset pricing models present serious problems in pricing a risk-free portfolio. Although they report high R^2 in the cross-section, this is typically achieved at the expense of estimated intercepts that are substantially greater than their theoretical values (i.e. the risk-free rate). To address this, we follow Nagel and Singleton (2011) and include the risk-free portfolio as an extra asset in our study.

We consider the orthogonal relationship between the risk-free rate and the pricing kernel for the following extra moment condition:^{*}

$$0 = \mathbb{E} \left[M(t+1) - \frac{1}{1 + R_f(t+1)} \right]. \quad (15.11)$$

In addition, we use the beta representation of Equation (15.9) to produce the unconditional expressions for the expected risk-free return and the expected return of asset i given by

$$\mathbb{E}[R_f(t+1)] = \frac{1}{\mathbb{E}[M(t+1)]} - 1, \quad (15.12)$$

$$\begin{aligned} \mathbb{E}[R_i(t+1)] &= \mathbb{E}[R_f(t+1)] - \mathbb{E}[1 + R_f(t+1)] \\ &\quad \times \text{Cov}[M(t+1), R_i(t+1)t], \end{aligned} \quad (15.13)$$

where in Equation (15.12) we are assuming that the risk-free asset is unconditionally orthogonal to $M(t+1)$ (i.e. the risk-free asset is a zero-beta asset). Equations (15.12) and (15.13) can be used to assess the fit of an estimated model to the cross section of average returns.

For the cross sectional tests we use the standard 25 portfolios of Fama and French (1993) ordered by size and book-to-market, in addition to the risk-free portfolio, therefore $N = 26$. These series as well as the SMB (small minus big) and HML (high minus low) factors from the FF three-factor model are available from Ken French's web page. The real per-capita consumption series was constructed using data from the Bureau of Economic Analysis. Specifically, we construct our quarterly series from nominal consumption of non-durables and services, seasonally adjusted, per capita (NIPA Table 7.1). Real consumption was calculated by deflating the nominal series by the PCE (personal consumption expenditures) price index, 2005 = 100 (NIPA Table 2.3.4).

Tables 15.10–15.12 present the results of the estimation by the Generalized Method of Moments (GMM) for the moment conditions in Equations (15.9) and (15.11).[†] We report the asymptotic t -stats, and the Wald and J_T tests based on the covariance matrices of pricing errors corrected for heteroscedasticity and serial correlation using the Newey–West estimator. We use the root of mean square errors (RMSE) to measure the fit of an estimated model to the cross-section of average returns. Figures 15.3–15.15 plot the fitted expected returns for the 26 portfolios against their realized average returns.[‡]

Table 15.10 and Figures 15.3–15.5 show the results for the unconditional CAPM, CCAPM and FF three-factor models. The J_T tests do not reject any conditional or unconditional model, so our inference

^{*} The grounds for this moment condition is that the risk-free rate is known at time t , therefore $1/[1 + R_f(t+1)] = \mathbb{E}[M(t+1) | \Omega(t)]$. By taking the unconditional mean on both sides yields equation (15.11).

[†] We use the identity weighting matrix for all estimates, based on the following reasons. First, we do not have theoretical arguments for giving more or less importance to a particular portfolio. Second, the number of moment conditions ($N = 26$) is large relative to our sample size ($T = 103$), so this choice avoids dealing with estimates that depend on unstable and near singular error covariance matrices.

[‡] The 25 portfolios, sorted by size and book-to-market ratio, are labelled with two digits. The first digit refers to the size quintile (1 indicating the smallest firms, 5 the largest), and the second digit refers to book-to-market quintile (1 indicating the portfolio with the lowest book-to-market ratio, 5 with the highest).

TABLE 15.10 Unconditional Asset Pricing Models, 1983Q2–2009Q4. Gmm Estimates of Pricing Kernel Coefficients for the Unconditional Models. The Models are Estimated Using Returns on the Fama and French (1993) Portfolios and A Risk-Free Portfolio ($N = 26$). The Identity-Weighting Matrix is Used in All Estimates. All Tests are Based on Covariance Matrices of Errors Corrected For Heteroscedasticity and Serial Correlation (Newey and West 1987). Lag Length in the Newey–West Estimator is $\text{floor}[4 \cdot (T/100)^{2/9}]$, Where $\text{floor}[x]$ Denotes the Integer Part of x . Asymptotic t -Stat in Parentheses. At The Bottom of the Table, We Report the P -Value For the J_t Test of the Null that all Pricing Errors are Zero. p -Value Wald *All* is the p -Value for the Asymptotic Wald Test that all Coefficients Except the Constant Are Zero. *Rmse* is the Root of Mean Square Errors And Measures the Fit of the Estimated Model to the Cross Section of Average Returns

	CAPM	CCAPM	FF Three-Factor
Constant	−0.027 (−0.66)	0.506 (1.31)	0.022 (0.60)
$R_m(t)$	−0.139 (−0.11)		
$\Delta c(t)$		−121.603 (−1.21)	
$R_m(t) - R_f(t)$			−2.058 (−1.25)
SMB(t)			0.259 (0.09)
HML(t)			−3.986 (−1.80)
p -Value J_T	0.44	0.43	0.32
p -Value Wald <i>All</i>	0.91	0.23	0.28
RMSE (%)	0.78	0.74	0.55

is based solely on the t -stats and Wald tests. The latter tests the hypothesis that all coefficients except the constant are zero. According to the Wald test, no model is significant at the 5% level; nevertheless, the FF three-factor model presents the best cross sectional adjustment ($RMSE = 0.55\%$).

Table 15.11 and Figures 15.6–15.10 present the results for the conditional CAPM models. A Wald test for the null hypothesis that the conditional CCAPM model does not improve the adjustment relative to the unconditional CCAPM model is included in the p -value Wald CAPM row (i.e. it tests whether the additional coefficients in the conditional CCAPM model are zero). The table shows that the CAPM conditional on Δf exhibits by far the best forecasting power for the cross-section of expected returns ($RMSE = 0.34\%$); moreover, in accordance with both Wald tests, it is the only significant one at the 5% level. The effect of the risk factor $R_m(t)$ is significant only through its interaction with $\Delta f(t - 4)$. This implies that a positive oil price shock accompanied by a subsequent decline in market return ($\Delta f(t - 4) \cdot R_m(t) < 0$) causes a drop in future portfolio returns. This effect strongly suggests that our variable also has forecasting power for returns on more disaggregated stock portfolios. Finally, CAPM models conditioned on $d - p$ and gap slightly outperform the FF three-factor model, but none of these models is statistically significant.

Table 15.12 and Figures 15.11–15.15 display the outcomes for conditional CCAPM models. Based on these outcomes, we note that the CCAPM conditional on Δf is the only one that outperforms the FF three-factor model; however, it is not significant at the 5% level. Also, the CCAPM model conditioned on R_f is statistically significant but is outperformed by several conditional and unconditional models ($RMSE = 0.60\%$).

In summary, the conditional CAPM on our variable has significant predictive power for the cross-section of expected stock returns and has a better fit than all unconditional and conditional models considered here.

TABLE 15.11 Conditional CCAPM Models, 1983Q2–2009Q4. GMM Estimates of Pricing Kernel Coefficients for the Conditional CCAPM Models on Variables are in the First Row. The Models are Estimated Using Returns on Fama and French's (1993) 25 Portfolios and the Risk-Free Asset ($N = 26$). The Identity-Weighting Matrix is Used in all Estimates. All Tests are Based on Covariance Matrices of Errors Corrected for Heteroscedasticity and Serial Correlation. Lag Length in the Newey–West Estimator is $\text{Floor}[4 \cdot (T/100)^{2/9}]$, Where $\text{Floor}[x]$ Denotes the Integer Part of x . Asymptotic T -Stat in Parentheses. At the Bottom of Table, We Report the p -Value for the J_t Test of the Null That All Pricing Errors are Zero. The p -Values Presented are for the Asymptotic Wald Tests. *All* Means All Coefficients Except the Constant. *CCAPM* Means the Coefficients that are not in the Unconditional CAPM Model. RMSE is the Root of Mean Square Errors and Measures the Adjustment of an Estimated Model to the Cross Section of Average Returns

	Δf	R_f	$d - p$	cay	gap
Constant	−1.262 (−2.61)	0.266 (0.41)	1.962 (0.38)	−0.413 (−0.75)	0.007 (0.09)
$\Delta f(t - 1)$	−2.744 (−2.18)				
$\Delta f(t - 2)$	−3.239 (−2.68)				
$\Delta f(t - 3)$	−1.387 (−0.46)				
$\Delta f(t - 4)$	6.932 (3.07)				
$R_f(t)$		−39.419 (−0.66)			
$d(t - 1) - p(t - 1)$			0.520 (0.38)		
$cay(t - 1)$				42.908 (1.13)	
$gap(t - 1)$					0.283 (0.03)
$R_m(t)$	1.157 (0.59)	−7.826 (−1.42)	−59.513 (−1.96)	1.602 (0.44)	−1.634 (−0.96)
$\Delta f(t - 1) \cdot R_m(t)$	−20.573 (−1.21)				
$\Delta f(t - 2) \cdot R_m(t)$	−17.374 (−1.63)				
$\Delta f(t - 3) \cdot R_m(t)$	11.929 (0.54)				
$\Delta f(t - 4) \cdot R_m(t)$	−63.751 (−2.92)				
$R_f(t) \cdot R_m(t)$		785.090 (1.42)			
$(d(t - 1) - p(t - 1)) \cdot R_m(t)$			−15.186 (−1.90)		
$cay(t - 1) \cdot R_m(t)$				−251.228 (−0.95)	
$gap(t - 1) \cdot R_m(t)$					89.443 (1.79)
p -Value J_T	0.07	0.32	0.33	0.33	0.35
p -Value Wald <i>All</i>	0.00	0.46	0.19	0.45	0.25
p -Value Wald <i>CAPM</i>	0.00	0.33	0.15	0.52	0.16
RMSE (%)	0.34	0.67	0.53	0.72	0.53

TABLE 15.12 Conditional CCAPM Models, 1983Q2–2009Q4. GMM Estimates of Pricing Kernel Coefficients for the Conditional CCAPM Models on Variables are in the First Row. The Models are Estimated Using Returns on Fama and French's (1993) 25 Portfolios and the Risk-Free Asset ($N = 26$). The Identity-Weighting Matrix is Used in all Estimates. All Tests are Based on Covariance Matrices of Errors Corrected for Heteroscedasticity and Serial Correlation. Lag Length In the Newey–West Estimator is $\text{Floor}[4 \cdot (T/100)^{2/9}]$, Where $\text{Floor}[X]$ Denotes the Integer part of x . Asymptotic t -Stat in Parentheses. At the Bottom of Table, We Report the p -Value for the J_t test of the Null That All Pricing Errors are Zero. The p -Values Presented are for the Asymptotic Wald Tests. *All* Means all Coefficients Except the Constant. *CCAPM* Means the Coefficients that are not in the Unconditional CCAPM Model. RMSE is the Root of Mean Square Errors and Measures the Adjustment of an Estimated Model to the Cross-Section of Average Returns

	Δf	R_f	$d - p$	cay	gap
Constant	0.811 (2.25)	1.033 (1.48)	1.877 (0.44)	0.375 (0.86)	0.485 (1.50)
$\Delta f(t - 1)$	−0.483 (−0.38)				
$\Delta f(t - 2)$	−1.858 (−2.15)				
$\Delta f(t - 3)$	0.384 (0.29)				
$\Delta f(t - 4)$	−0.407 (−0.17)				
$R_f(t)$		−186.224 (−2.18)			
$d(t - 1) - p(t - 1)$			0.364 (0.32)		
$cay(t - 1)$				19.948 (0.76)	
$gap(t - 1)$					−0.253 (−0.04)
$\Delta c(t)$	−308.738 (−2.77)	−320.549 (−1.83)	−390.098 (−0.41)	−160.464 (−1.57)	−144.812 (−1.51)
$\Delta f(t - 1) \cdot \Delta c(t)$	116.051 (0.35)				
$\Delta f(t - 2) \cdot \Delta c(t)$	−134.842 (−0.59)				
$\Delta f(t - 3) \cdot \Delta c(t)$	343.002 (1.09)				
$\Delta f(t - 4) \cdot \Delta c(t)$	651.304 (1.46)				
$R_f(t) \cdot \Delta c(t)$		34600.413 (2.88)			
$(d(t - 1) - p(t - 1)) \cdot \Delta c(t)$			−68.620 (−0.27)		
$cay(t - 1) \cdot \Delta c(t)$				1269.334 (0.34)	
$gap(t - 1) \cdot \Delta c(t)$					1302.922 (1.05)
p -Value J_T	0.08	0.34	0.33	0.32	0.34
p -Value Wald <i>All</i>	0.10	0.03	0.44	0.33	0.47
p -Value Wald <i>CCAPM</i>	0.24	0.01	0.93	0.74	0.54
RMSE (%)	0.48	0.60	0.73	0.69	0.68




FIGURE 15.3 Realized vs. expected returns by CAPM.




FIGURE 15.4 Realized vs. expected returns by CCAPM.




FIGURE 15.5 Realized vs. expected returns by FF three-factor.




FIGURE 15.6 Realized vs. expected returns by CAPM conditioned on Δf .




FIGURE 15.7 Realized vs. expected returns by CAPM conditioned on R_f .




FIGURE 15.8 Realized vs. expected returns by CAPM conditioned on $d - p$.




FIGURE 15.9 Realized vs. expected returns by CAPM conditioned on *cay*.




FIGURE 15.10 Realized vs. expected returns by CAPM conditioned on *gap*.




FIGURE 15.11 Realized vs. expected returns by CCAPM conditioned on Δf .




FIGURE 15.12 Realized vs. expected returns by CCAPM conditioned on R_f .




FIGURE 15.13 Realized vs. expected returns by CCAPM conditioned on $d - p$.




FIGURE 15.14 Realized vs. expected returns by CCAPM conditioned on cay .




FIGURE 15.15 Realized vs. expected returns by CCAPM conditioned on *gap*.

15.7 Conclusions

Although the predictability of stock market excess returns has been associated with the business cycle, the evidence that supports this relationship is far from being conclusive. Moreover, when the sample is extended to include the period of the sub-prime crisis, none of the popular predictors exhibit forecasting power for market excess returns. In this paper, we show that such a relationship exists.

We find that unexpected oil price changes, a non-persistent variable with deep macroeconomic roots, have significant forecasting power for stock returns at short horizons. Our variable, proxied by futures returns on crude oil, shows statistically and economically significant predictive power for stock returns at horizons from one to three quarters. Its predictive power outperforms those of the risk-free rate, the dividend–price ratio, the consumption–wealth ratio and the output gap, with quarterly $\bar{R}^2$ between 6% and 7%. This result is robust against the inclusion of other variables and out-of-sample tests. However, at longer horizons, none of the variables displays significant forecasting ability. Our results also validate the recent findings of Boudoukh *et al.* (2008) that unstable results in previous studies in this literature are due to the high persistence of the predictors used.

Our variable also shows a significant forecasting power for the cross section of expected returns. We build a conditional CAPM model on oil price shocks, which shows high statistical significance and better adjustment than all conditional and unconditional models considered, including the Fama and French (1993) three-factor model.

From a practical perspective, unlike variables based on macroeconomic series, such as the consumption–wealth ratio and the output gap, our variable can be directly observed and is available on a daily basis at no cost. These characteristics make use of our variable by potential investors highly feasible.

Finally, an open question motivated by the emerging literature on the financialization of commodities is how well is the forecasting ability of other commodities on stock returns. The correlation of indexed commodities and oil has increased dramatically since 2004 because of the speculative trading in futures markets, suggesting that these commodities would also have some predictive power. We leave this topic for future research.

Acknowledgements

We would like to thank Augusto Castillo, Gonzalo Cortazar, Rodrigo Fuentes, Borja Larrain, Chris Telmer, Jose Tessada, Eduardo Walker and seminar participants at UC. Any errors or omissions are the responsibility of the authors. Casassus acknowledges financial support from FONDECYT (grant No. 1110841) and from Grupo Security through FinanceUC. Higuera acknowledges financial support from CONICYT and FinanceUC.

References

- Acharya, V.V., Lochstoer, L.A. and Ramadorai, T., Limits to arbitrage and hedging: Evidence from commodity markets. NBER Working Paper No. 16875, 2011.
- Aguiar-Conraria, L. and Wen, Y., Understanding the large negative impact of oil shocks. *J. Money, Credit Bank.*, 2007, **39**(4), 925–944.
- Ang, A. and Bekaert, G., Stock return predictability: Is it there? *Rev. Financ. Stud.*, 2007, **20**(3), 651–707.
- Apergis, N. and Miller, S.M., Do structural oil-market shocks affect stock prices? *Energy Econ.*, 2009, **31**(4), 569–575.
- Bachmeier, L., Monetary policy and the transmission of oil shocks. *J. Macroecon.*, 2008, **30**(4), 1738–1755.
- Bakshi, G., Panayotov, G. and Skoulakis, G., The Baltic Dry Index as a predictor of global stock returns, commodity returns, and global economic activity. Working Paper, University of Maryland, 2011.
- Balke, N.S., Brown, S.P.A. and Yucel, M.K., Oil price shocks and the U.S. economy: Where does the asymmetry originate? *Energy J.*, 2002, **23**(3), 27–52.
- Barsky, R.B. and Kilian, L., Oil and the macroeconomy since the 1970s. *J. Econ. Perspect.*, 2004, **18**(4), 115–134.
- Basak, S. and Pavlova, A., A model of financialization of commodities. Working paper, London Business School, 2012.
- Berkowitz, J. and Kilian, L., Recent developments in bootstrapping time series. *Econometr. Rev.*, 2000, **19**(1), 1–48.
- Bernanke, B.S., Irreversibility, uncertainty, and cyclical investment. *Q. J. Econ.*, 1983, **98**(1), 85–106.
- Bernanke, B.S., Gertler, M. and Watson, M.W., Systematic monetary policy and the effects of oil price shocks. *Brook. Pap. Econ. Activ.*, 1997, **1997**(1), 91–157.
- Boudoukh, J., Richardson, M. and Whitelaw, R.F., The myth of long-horizon predictability. *Rev. Financ. Stud.*, 2008, **21**(4), 1577–1605.
- Breedon, D.T. and Litzenberger, R.H., Prices of state-contingent claims implicit in option prices. *J. Business*, 1978, **51**(4), 621–651.
- Buyuksahin, B. and Robe, M.A., Speculators, commodities and cross-market linkages. Working paper, American University, 2010.
- Campbell, J.Y., Stock returns and the term structure. *J. Financ. Econ.*, 1987, **18**(2), 373–399.
- Campbell, J.Y., A variance decomposition for stock returns. *Econ. J.*, 1991, **101**(405), 157–179.
- Campbell, J.Y. and Thompson, S.B., Predicting excess stock returns out of sample: Can anything beat the historical average? *Rev. Financ. Stud.*, 2008, **21**(4), 1509–1531.
- Carruth, A.A., Hooker, M.A. and Oswald, A.J., Unemployment equilibria and input prices: Theory and evidence from the United States. *Rev. Econ. Statist.*, 1998, **80**(4), 621–628.
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**(5), 2283–2332.
- Ciner, C., Energy shocks and financial markets: Nonlinear linkages. *Stud. Nonlin. Dynam. Econometr.*, 2001, **5**(3), 203–212.
- Clark, T.E. and McCracken, M.W., Tests of equal forecast accuracy and encompassing for nested models. *J. Econometr.*, 2001, **105**, 85–110.

- Clark, T.E. and McCracken, M.W., Evaluating direct multistep forecasts. *Econometr. Rev.*, 2005, **24**(4), 369–404.
- Cochrane, J.H., *Asset Pricing*, 2005 (Princeton University Press: Princeton).
- Cologni, A. and Manera, M., Oil prices, inflation and interest rates in a structural cointegrated VAR model for the G-7 countries. *Energy Econ.*, 2008, **30**(3), 856–888.
- Cologni, A. and Manera, M., The asymmetric effects of oil shocks on output growth: A Markov-switching analysis for the G-7 countries. *Econ. Model.*, 2009, **26**(1), 1–29.
- Cooper, I. and Priestley, R., Time-varying risk premiums and the output gap. *Rev. Financ. Stud.*, 2009, **22**(7), 2801–2833.
- Davis, S.J. and Haltiwanger, J., Sectoral job creation and destruction responses to oil price changes. *J. Monet. Econ.*, 2001, **48**(3), 465–512.
- Diebold, F.X. and Mariano, R.S., Comparing predictive accuracy. *J. Business Econ. Statist.*, 1995, **13**(3), 253–263.
- Driesprong, G., Jacobsen, B. and Maat, B., Striking oil: Another puzzle? *J. Financ. Econ.*, 2008, **89**(2), 307–327.
- Etula, E., Broker{dealer risk appetite and commodity returns. Working paper, Federal Reserve Bank of New York, 2010.
- Fama, E.F. and French, K.R., Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *J. Business*, 1987, **60**(1), 55–73.
- Fama, E.F. and French, K.R., Dividend yields and expected stock returns. *J. Financ. Econ.*, 1988, **22**(1), 3–25.
- Fama, E.F. and French, K.R., Common risk factors in the returns on stocks and bonds. *J. Financ. Econ.*, 1993, **33**(1), 3–56.
- Ferderer, J.P., Oil price volatility and the macroeconomy. *J. Macroecon.*, 1996, **18**(1), 1–26.
- Finn, M.G., Perfect competition and the effects of energy price increases on economic activity. *J. Money, Credit Bank.*, 2000, **32**(3), 400–416.
- Gronwald, M., Large oil shocks and the US economy: Infrequent incidents with large effects. *Energy J.*, 2008, **29**(1), 151–171.
- Hamilton, J.D., Oil and the macroeconomy since World War II. *J. Polit. Econ.*, 1983, **91**(2), 228–248.
- Hamilton, J.D., A neoclassical model of unemployment and the business cycle. *J. Polit. Econ.*, 1988, **96**(3), 593–617.
- Hamilton, J.D., This is what happened to the oil price–macroeconomy relationship. *J. Monet. Econ.*, 1996, **38**(2), 215–220.
- Hamilton, J.D., What is an oil shock?. *J. Econometr.*, 2003, **113**(2), 363–398.
- Hamilton, J.D., Oil and the macroeconomy. In *The New Palgrave Dictionary of Economics*, 2nd ed., edited by S.N. Durlauf and L.E. Blume, 2008 (Palgrave Macmillan: Basingstoke). Available online at: http://www.dictionaryofeconomics.com/article?id=pde2008_E000233
- Hamilton, J.D. and Herrera, A.M., Comment: Oil shocks and aggregate macroeconomic behavior: The role of monetary policy. *J. Money, Credit Bank.*, 2004, **36**(2), 265–286.
- Harvey, D.I., Leybourne, S.J. and Newbold, P., Tests for forecast encompassing. *J. Business Econ. Statist.*, 1998, **16**(2), 254–259.
- Hooker, M.A., What happened to the oil price–macroeconomy relationship? *J. Monet. Econ.*, 1996, **38**(2), 195–213.
- Huang, R.D., Masulis, R.W. and Stoll, H.R., Energy shocks and financial markets. *J. Fut. Mkts*, 1996, **16**(1), 1–27.
- Huntington, H.G., Oil shocks and real U.S. income. *Energy J.*, 2007, **28**(4), 31–46.
- Inoue, A. and Kilian, L., In-sample or out-of-sample tests of predictability: Which one should we use? *Econometr. Rev.*, 2005, **23**(4), 371–402.
- Jones, C. and Kaul, G., Oil and the stock markets. *J. Finance*, 1996, **51**(2), 463–491.

- Kilian, L., Small-sample confidence intervals for impulse response functions. *Rev. Econ. Statist.*, 1998, **80**(2), 218–230.
- Kilian, L., Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions? *J. Appl. Econometr.*, 1999, **14**(5), 491–510.
- Kilian, L., A comparison of the effects of exogenous oil supply shocks on output and inflation in the G7 countries. *J. Eur. Econ. Assoc.*, 2008, **6**(1), 78–121.
- Kilian, L., Not All oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *Am. Econ. Rev.*, 2009, **99**(3), 1053–1069.
- Kilian, L. and Park, C., The impact of oil price shocks on the U.S. stock market. *Int. Econ. Rev.*, 2009, **50**(4), 1267–1287.
- Kliesen, K.L., Oil and the U.S. macroeconomy: An update and a simple forecasting exercise. *Fed. Reserve Bank St. Louis Rev.*, 2008, **90**(5), 505–516.
- Kumar, M.S., The forecasting accuracy of crude oil futures prices. *Staff Papers – International Monetary Fund*, 1992, **39**(2), 432–461.
- Leduc, S. and Sill, K., A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns. *J. Monet. Econ.*, 2004, **51**(4), 781–808.
- Lee, K. and Ni, S., On the dynamic effects of oil price shocks: A study using industry level data. *J. Monet. Econ.*, 2002, **49**(4), 823–852.
- Lee, K., Ni, S. and Ratti, R.A., Oil shocks and the macroeconomy: The role of price variability. *Energy J.*, 1995, **16**(4), 39–56.
- Lettau, M. and Ludvigson, S.C., Consumption, aggregate wealth, and expected stock returns. *J. Finance*, 2001a, **56**(3), 815–849.
- Lettau, M. and Ludvigson, S.C., Resurrecting the (C)CAPM: A cross sectional test when risk premia are time-varying. *J. Polit. Econ.*, 2001b, **109**(6), 1238–1287.
- Lewellen, J., Nagel, S. and Shanken, J., A skeptical appraisal of asset pricing tests. *J. Financ. Econ.*, 2010, **96**(2), 175–194.
- Lilien, D.M., Sectoral shifts and cyclical unemployment. *J. Polit. Econ.*, 1982, **90**(4), 777–793.
- Lustig, H.N. and Van Nieuwerburgh, S.G., Housing collateral, consumption insurance, and risk premia: An empirical perspective. *J. Finance*, 2005, **60**(3), 1167–1219.
- Ma, C.W., Forecasting efficiency of energy futures prices. *J. Fut. Mkts*, 1989, **9**(5), 393–419.
- Masters, M.W., Testimony before Committee on Homeland Security and Governmental Affairs of the United States Senate, May 2008.
- McCracken, M.W., Asymptotics for out of sample tests of Granger causality. *J. Econometr.*, 2007, **140**, 719–752.
- Mork, K.A., Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *J. Polit. Econ.*, 1989, **97**(3), 740–744.
- Nagel, S. and Singleton, K.J., Estimation and evaluation of conditional asset pricing models. *J. Finance*, 2011, **66**(3), 873–909.
- Nandha, M. and Faff, R., Does oil move equity prices? A global view. *Energy Econ.*, 2008, **30**(3), 986–997.
- Newey, W.K. and West, K.D., A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 1987, **55**(3), 703–708.
- Newey, W.K. and West, K.D., Automatic lag selection in covariance matrix estimation. *Rev. Econ. Stud.*, 1994, **61**(4), 631–653.
- Park, J. and Ratti, R.A., Oil price shocks and stock markets in the U.S. and 13 European countries. *Energy Econ.*, 2008, **30**(5), 2587–2608.
- Sadorsky, P., Oil price shocks and stock market activity. *Energy Econ.*, 1999, **21**(5), 449–469.
- Santos, T. and Veronesi, P., Labor income and predictable stock returns. *Rev. Financ. Stud.*, 2006, **19**(1), 1–44.
- Sharpe, W.F., Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance*, 1964, **19**(3), 425–442.

- Sims, C.A., Macroeconomics and reality. *Econometrica*, 1980, **48**(1), 1–48.
- Singleton, K.J., Investor owns and the 2008 boom/bust in oil prices. Working paper, Stanford University, 2012.
- Stambaugh, R.F., Predictive regressions. *J. Financ. Econ.*, 1999, **54**(3), 375–421.
- Stine, R.A., Estimating properties of autoregressive forecasts. *J. Am. Statist. Assoc.*, 1987, **82**(400), 1072–1078.
- Switzer, L.N. and El-Khoury, M., Extreme volatility, speculative efficiency, and the hedging effectiveness of the oil futures markets. *J. Fut. Mkts*, 2007, **27**(1), 61–84.
- Tang, K. and Xiong, W., Index investment and the financialization of commodities. *Financ. Anal. J.*, 2012, **68**(6), 54–74.
- Valkanov, R., Long-horizon regressions: Theoretical results and applications. *J. Financ. Econ.*, 2003, **68**(2), 201–232.
- Wei, C., Energy, the stock market, and the Putty-Clay Investment model. *Am. Econ. Rev.*, 2003, **93**(1), 311–323.
- Welch, I. and Goyal, A., A comprehensive look at the empirical performance of equity premium prediction. *Rev. Financ. Stud.*, 2008, **21**(4), 1455–1508.

Appendix A: Tests for Out-of-Sample Predictability

This appendix presents the three metrics we use to test the out-of-sample performance of the predictors. First, we estimate the forecast errors for the benchmark and competing models in Equations (15.4) and (15.5) as

$$\text{benchmark : } \hat{u}_1(t+1) = [R_m(t+1) - R_f(t+1)] - \hat{\alpha}_1(t), \quad (\text{A.1})$$

$$\begin{aligned} \text{competing : } \hat{u}_2(t+1) = & [R_m(t+1) - R_f(t+1)] \\ & - \hat{\alpha}_2(t) - \hat{\beta}(t)'X(t), \end{aligned} \quad (\text{A.2})$$

for $t = Q, \dots, T-1$, and the coefficients $\hat{\alpha}_1(t)$, $\hat{\alpha}_2(t)$ and $\hat{\beta}(t)$ are estimated with data through periods $1, \dots, t$. Then, one-step-ahead forecasts from the competing model can be compared with forecasts from the benchmark model (that is, a restricted version of the competing model) by using statistics based on the time series $\hat{u}_1(t+1)$ and $\hat{u}_2(t+1)$.

The first test we use for out-of-sample predictability is the forecast encompassing test of Clark and McCracken (2001). To clarify how it works, we follow Harvey *et al.* (1998) and specify a regression of the excess stock return on a weighted average of forecasted values from the benchmark and competing models:

$$\begin{aligned} & R_m(t+1) - R_f(t+1) \\ & = (1-\lambda)[\alpha_1] + \lambda[\alpha_2 + \beta'X(t)] + v(t+1), \end{aligned} \quad (\text{A.3})$$

where $0 \leq \lambda \leq 1$ and $v(t+1)$ is an error term. Substituting both forecasts from Equations (15.4) and (15.5) yields

$$u_1(t+1) = \lambda[u_1(t+1) - u_2(t+1)] + v(t+1). \quad (\text{A.4})$$

Then, as λ is also the coefficient of the regression model in equation (A.4):

$$\lambda = \frac{\text{Cov}[u_1(t+1), u_1(t+1) - u_2(t+1)]}{\text{Var}[u_1(t+1) - u_2(t+1)]}. \quad (\text{A.5})$$

Thus, the combined forecast will have a smaller expected squared error than the benchmark model forecast unless the covariance between $u_1(t+1)$ and $u_1(t+1) - u_2(t+1)$ is zero (i.e. $\lambda = 0$). This way, Clark and McCracken (2001) tests the null hypothesis that $\lambda \leq 0$ and is given by

$$\text{ENC-NEW} = \frac{P \sum_{t=Q}^{T-1} (\hat{u}_1(t+1)^2 - \hat{u}_1(t+1)\hat{u}_2(t+1))}{\sum_{t=Q}^{T-1} \hat{u}_2(t+1)^2} \quad (\text{A.6})$$

Under the null hypothesis that the benchmark model encompasses the competing model, the covariance between series $u_1(t+1)$ and $u_1(t+1) - u_2(t+1)$ will be less than or equal to zero. Under the alternative that the competing model contains added information, the covariance should be positive. Hence, the encompassing test presented above is one-sided. Clark and McCracken (2001) demonstrate that the limiting distribution of ENC-NEW is not normal when the forecasts are nested under the null, but they provide asymptotic critical values for this statistic.

The second test used is the one developed by McCracken (2007). This test, unlike the one proposed by Diebold and Mariano (1995) in the context of non-nested models, allows for comparison of predictive accuracy between nested models. In particular, we use it to test for equality of the mean squared forecasting errors (MSE) from the benchmark and competing models, which is given by

$$\begin{aligned} \text{MSE-F} &= \frac{P \sum_{t=Q}^{T-1} (\hat{u}_1(t+1)^2 - \hat{u}_2(t+1)^2)}{\sum_{t=Q}^{T-1} \hat{u}_2(t+1)^2}, \\ &= P \left[\frac{\text{MSE}_1 - \text{MSE}_2}{\text{MSE}_2} \right] \end{aligned} \quad (\text{A.7})$$

where $\text{MSE}_j = \sum_{t=Q}^{T-1} \hat{u}_j(t+1)^2 / P$, $j = 1, 2$. Based upon the value of this statistic the null of equal MSE is either rejected or not rejected. McCracken (2007) shows that when the two models are nested the alternative is one-sided, rather than two-sided. Moreover, since the asymptotic distribution of MSE-F under the null is non-standard, tables of asymptotically valid critical values are provided by McCracken (2007).

Clark and McCracken (2001) use simulations to examine the small-sample properties of the ENC-NEW and MSE-F tests. They report that although both tests have good sample size properties, the ENC-NEW test is clearly the more powerful out-of-sample test of predictive ability. While this evidence indicates that the inference from the ENC-NEW test is more reliable, Welch and Goyal (2008) highlight an important problem of encompassing tests in general. The ENC-NEW test uses the entire out-of-sample test to estimate the parameter λ , but an investor trying to use a combined forecast to predict $R_m(t+1) - R_f(t+1)$ will only have the information available up to t to calculate the combination coefficient λ .

The final measure of out-of-sample forecasting performance is the out-of-sample R^2 , R_{OS}^2 . This statistic is the analogue to the in-sample R^2 and in terms of our notation is computed as

$$\begin{aligned} R_{\text{OS}}^2 &= 1 - \frac{\sum_{t=Q}^{T-1} \hat{u}_2(t+1)^2}{\sum_{t=Q}^{T-1} \hat{u}_1(t+1)^2} = \frac{\text{MSE}_1 - \text{MSE}_2}{\text{MSE}_1} \\ &= \frac{\text{MSE-F}}{P} \left(\frac{\text{MSE}_2}{\text{MSE}_1} \right). \end{aligned} \quad (\text{A.8})$$

As can be seen from Equation (A.8), if R_{OS}^2 is positive then the competing model has a lower MSE than the benchmark model. Also, as shown in the last equality, R_{OS}^2 is not a statistic that provides new

information with respect to the other tests, since it is merely a scaled-up version of the MSE-F statistic.* That is, predictor variables with greater MSE-F will also exhibit greater R_{OS}^2 . Then, this could also be considered a test of equal MSE, assuming that it has an asymptotic distribution.†

As mentioned above, in the context of one-step-ahead forecasts, Clark and McCracken (2001) and McCracken (2007) provide asymptotic critical values for the ENC-NEW and MSE-F statistics, respectively. These critical values depend on two parameters: $\pi = P/Q$ and $K_2 - 1$, the number of variables included in $X(t)$. Because the tables with the critical values for these non-standard tests do not contain the particular value of π chosen by us, we follow Clark and McCracken (2005) and obtain these values with an inference technique based on bootstrapping. In addition, based on the bootstrapped time series, we obtain the empirical distribution of the R_{OS}^2 statistic and critical values for the tests. In particular, we use a parametric bootstrap Berkowitz and Kilian (2000) and our algorithm has five steps, which we briefly describe below:‡

1. We estimate a bivariate VAR for the excess stock return, $R_m(t) - R_f(t)$, and the variable $X(t)$ under the null hypothesis of non-predictability. The model is estimated with OLS and using the full sample. The excess return is modelled according to Equation (15.4) and for the variable $X(t)$ the optimal number of lags of $R_m(t) - R_f(t)$ and $X(t)$ were chosen with the AIC criterion.
2. The coefficients of the VAR were adjusted for the small-sample bias using the procedure of Kilian (1998) with 10 000 bootstrap draws.
3. We bootstrapped 999 time series for the excess stock return and the variable $X(t)$ by drawing from the rescaled sample residuals with replacement (Berkowitz and Kilian 2000), using the adjusted VAR coefficients and initial observations selected by sampling from actual data (Stine 1987).
4. Each artificial bivariate time series is used to estimate the benchmark and competing models (Equations (15.4) and (15.5)) in a recursive way. Forecasting errors are calculated according to Equations (A.1) and (A.2) and using the sample portions described above. The ENC-NEW, MSE-F and R_{OS}^2 statistics were calculated based on these estimated forecasting errors with Equations (A.6), (A.7) and (A.8).
5. For each statistic, critical values are simply computed as percentiles of the corresponding empirical distribution. The p -values are calculated using the standard method.

* The MSE-F, as derived originally by McCracken (2007), is designed to be used with any loss function. Here we use only one particular case.

† Just as the in-sample R^2 has an adjusted counterpart for degrees of freedom, $\bar{R}^2$, Welch and Goyal (2008) use a version of the R_{OS}^2 adjusted for degrees of freedom. However, since the forecasting errors are not part of the OLS estimation, and therefore there is no loss of degrees of freedom in its calculation, we consider it inappropriate to adjust this statistic.

‡ For more details on this methodology, see Clark and McCracken (2005).

16

Time-frequency Analysis of Crude Oil and S&P 500 Futures Contracts

16.1	Introduction	337
16.2	Literature Review	339
16.3	Data	341
16.4	Methodology	345
	Cross-Spectral Analysis • Wavelets • Cross Wavelets • Cross Wavelet Phase Angles	
16.5	Empirical Results	348
	Cospectral Densities • Wavelets	
16.6	Conclusions	354
	Acknowledgements	355
	Software Acknowledgements	356
	References	356

Joseph McCarthy

Alexei G. Orlov

We use frequency-domain techniques, namely wavelets and cross-spectra, to examine the association between the daily prices of crude oil futures and daily S&P 500 futures closing prices over the past several decades. We investigate contemporaneous and lag-lead relationships in levels and returns. It is our belief that the wavelet and cross-spectral analyses employed in this paper offer insights regarding the relationship between oil prices and stock returns that are not apparent from a conventional time-domain framework. Our findings cast doubt on the purported negative relationship between oil and the U.S. stock market. Our analysis suggests that oil prices lead oil volume, and S&P 500 trading volume leads S&P 500 prices.

Keywords: Wavelets in finance; Comovements; Financial futures; Commodity prices; Cross-spectrum in finance

JEL Classification: C1, C14, G1, G10, G13

16.1 Introduction

We use frequency-domain techniques to study the effects of oil price movements on U.S. stock returns. We also seek to help explain the dynamics of the stock market by examining levels and returns in the crude oil market. Specifically, we study the linkages between volatility in the futures price of crude oil and the futures price of S&P 500 contracts, as well as the relationship between the volume of crude oil futures and the volume of the S&P 500 futures contracts. We also investigate the lead-lag relationship between crude oil and stock prices, stock returns and trading volumes.

Conventional wisdom (as delineated by, for example, Sauter and Awerbuch (2003)) suggests that changes in oil prices should be negatively related to macroeconomic and financial indicators. Higher oil prices may depress economic activity, thereby reducing personal income and wealth. Incomes may become lower due to higher unemployment and lower wages, while wealth may be negatively affected through reduced stock values of companies held by individuals.

There are several channels through which oil prices can affect financial volatility. First, capital market theory predicts that higher future costs associated with increased oil prices would translate into higher risk of holding the affected assets. Since the increased risk is bound to manifest in higher volatility, one would expect higher oil prices to be reflected in higher financial volatility and instability. Second, if there is a (negative) relationship between oil prices and the stock market in levels, this relationship is likely to be carried on to the second moments. It should be noted that volatility of oil prices *per se* would be indicative of a riskier economic environment, thereby reinforcing the predicted positive link between volatilities of oil prices and financial indicators.

Finally, one can contend that oil price movements and financial performance find their highest correlation at a lag or a lead: for example, the expected future higher price of oil may negatively affect the current value of assets and positively affect current financial volatility. We use cross-spectral and wavelet analyses to answer empirically the following questions stemming from the theoretical premises outlined above.

H1: Is there a negative (or positive) contemporaneous correspondence between oil and stocks, either in levels or returns?

H2: Is there a lead-lag relationship among trading volumes and prices in the oil and stock markets? If so, what is the lag or lead time that produces the strongest correspondence?

H3: Is there a negative (or positive) correspondence between the (levels of) oil prices and stock returns at a lag or a lead? If so, what is the lag or lead time that produces the strongest correspondence?

One of the first comprehensive studies on the relationship between oil prices and macroeconomic activity was Hamilton (1983), who found that higher oil prices led to slower economic growth in the U.S. for the period 1948–1980. Hamilton's conclusion was not entirely surprising, as his time span included two major oil shocks. Subsequent studies confirmed the negative relationship between oil prices and U.S. economic activity for more recent periods, as well as for other oil-consuming countries (e.g. Hooker (1996), Rotemberg and Woodford (1996), Hamilton (2000) and Yang *et al.* (2002)). These studies report a fairly substantial decrease in economic activity due to higher oil prices: a 10% oil price increase results in a decline in GDP growth of 0.6 to 2.5% (Sauter and Awerbuch 2003).*

In contrast to the established negative relationship between oil prices and the economy in the 1970s and 1980s, more recently lower prices of oil and natural gas did not appear to help propel the U.S. and other economies out of recession. Our ability to unequivocally declare the breakdown of negative correlations between oil prices and economic and financial performance is, of course, impaired by the complications of the ongoing financial crisis with its various causes and effects. It is thus imperative to revisit the empirical analysis of oil prices using more extensive data sets as well as frequency domain techniques that can help refine the conventional time-domain analyses.

It should be pointed out that prior to the mid-1980s, most oil price changes were price increases, whereas since the late 1980s we started to observe both positive and negative changes in oil prices (Sauter and Awerbuch 2003). This increased oil price volatility should prompt a re-examination of some of the earlier studies, as their conclusions may not pertain to the most recent patterns in the data. In fact, Mork (1989), Mork *et al.* (1994), Lee *et al.* (1995) and Lee and Zeng (2011), among others, report the asymmetric effects of oil prices for U.S. and other industrialized economies: the impact

* Some researchers (e.g. Ciner (2001)) find an even stronger relationship between oil prices and the stock market in the 1990s relative to the 1970s or 1980s.

on economic activity of oil price increases can be quite negative, while the impact of price decreases was insignificant.*

It should be made clear that the focus of our paper is rather narrow: we study the relationship between oil and security prices, as well as the relationship among trading volumes and prices. We do not purport to explain the state of the macroeconomy (i.e. unemployment, GDP, etc.). Instead, we use movements in S&P 500 futures contracts to proxy for financial volatility and investigate the link between said volatility and the oil market.

The subject of this paper is important in yet another respect. The debate surrounding fossil-based technologies *vis-a-vis* green energy is ongoing and impacts a large amount of investment capital. Although the primary focus of this debate is on climate change and the sustainability of natural resources, the effect of old versus new technologies on the state of national economies also plays a non-trivial role in this debate. It is thus interesting to determine if oil price fluctuations have disruptive effects on equity markets. If so, there is an additional argument to be made in favour of fostering the development and refinement of non-fossil energy sources. If, on the other hand, there are no such effects found in the most recent data, one cannot comfortably offer the financial stability argument when advocating the switch to alternative energy sources.

The objective of this paper is to use frequency domain techniques, such as wavelets and cross-spectra, to examine the association between the daily prices of crude oil futures and daily S&P 500 futures closing prices over the past several decades. We believe that the wavelet and cross-spectral analyses employed in this paper offer insights regarding the relationship between oil prices and financial indicators that are not apparent from a conventional time-domain framework. We also believe that our research introduces a useful way of looking at an important question: Are the recent fluctuations in the oil markets associated with greater comovements of oil prices and financial indicators? Our research allows one to judge to what extent the relationship between oil and the stock market has changed over time and if there is a lag-lead effect of (expected) oil price changes on the financial stability of the U.S. economy.†

Our extensive data sample includes a wide range of price movements, both with respect to levels, as well as rates and direction of change. We find a positive correspondence between oil prices and S&P 500 in levels, and our findings cast doubt on the purported negative relationship between oil and the U.S. stock market. Also, our analysis suggests that oil prices lead oil volume, while S&P 500 trading volume leads S&P 500 prices. These findings persist over a large number of time scales and across a wide range of Fourier frequencies.

16.2 Literature Review

Most empirical studies that focus on the role of oil prices in the economy, starting with Hamilton's (1983) seminal paper and continuing with Hooker (1996), Rotemberg and Woodford (1996), Hamilton (2000) and others, report significant negative effects of oil prices on macroeconomic activity. One approach to measure the effect of oil prices on financial indicators is through estimating a beta for oil. Recall that an asset's beta measures the covariance between that asset's value and the value of a diversified portfolio or a market index. Thus, an estimated negative beta for oil would indicate a negative covariance risk and, therefore, a negative relationship between oil prices and the stock market. Awerbuch (1993), Awerbuch and Deehan (1995) and Bolinger *et al.* (2002) use this approach and report a negative beta for oil and

* Ferderer (1996) suggests that the asymmetric effects of oil prices are due to (i) sectorial shifts of specialized labour and capital, and (ii) the uncertainty that may induce companies to postpone their investment. The latter issue of the impact of oil price uncertainty on investment was further investigated by Elder and Serletis (2010).

† Only recently have researchers started to rigorously investigate the possibility that oil prices may have predictive power in forecasting stock markets (i.e. Driesprong *et al.* (2008)).

natural gas. Such a result is indicative of a possible 'double whammy effect' (Sauter and Awerbuch 2003) of the higher oil prices: consumers feel the pinch at the pump and when they pay their utility bills, and they also see the decline in their wealth due to the lower value of their assets.

Many of the oil price increases, particularly after 1986, have been followed by large decreases (Hamilton 1996). Thus oil price volatility is likely to have played an increasingly important role in the last two decades. Lee *et al.* (1995) emphasize that oil price volatility may be an important variable in accounting for economic fluctuations. Including both the oil price volatility measure and the magnitude of an oil price shock (relative to the trend) in their analysis, Lee *et al.* (1995) conclude that an oil price shock has a smaller impact on the macroeconomy in a high volatility environment, i.e. when such a shock is less of a surprise. Ferderer (1996) reports that output is negatively affected by both levels and volatility of oil prices, with volatility having a more pronounced effect, and points to a possible lag between oil price changes and macroeconomic activity of about a year. A similar—and more assertive—conclusion is reached by Hooker (1996): oil price levels do not have predictive power over output, while volatility of oil prices does.

The possibility of a nonlinear relationship between oil prices and the stock markets has been advocated by Ciner (2001) and Reboredo (2010), among others. In particular—and most relevant to our study—Ciner finds a strong bidirectional causality between crude oil futures returns and the S&P 500 index returns. The interdependence of the stock and oil futures prices has recently been confirmed by Vo (2011). Thus, there appears to be a nonlinear feedback effect between oil and the stock market. The semiparametric and non-parametric methods used in this paper do not confine us to a set of linear relationships and are flexible enough to find non-linear links between oil prices and stock returns.

Although macroeconomic activity is closely related to financial indicators and, in particular, stock market performance, it must be acknowledged that the value of the stock market is not always synonymous with the current state of the economy. Huang *et al.* (1996) and Jones and Kaul (1996) were among the first comprehensive studies that examined the interplay between oil prices and changes in stock market fluctuations. Jones and Kaul (1996) report that oil prices help predict stock market returns in the U.S., Canada and Japan, and that the effect of oil prices on stock markets is negative both contemporaneously and at a lag. The results suggest that oil shocks affect the stock markets through their effect on current and future real cash flows, which is consistent with the cash-flow dividend valuation model (Campbell 1991).

Huang *et al.* (1996) study the link between daily oil futures returns and daily U.S. stock returns using vector autoregressions (VARs). Unlike Jones and Kaul (1996), they did not find any evidence that oil futures returns have an impact on stock market indices such as S&P 500. However, the findings of Jones and Kaul are confirmed by Sadorsky (1999), who applies VARs to monthly data and reports a negative and significant impact of an oil shock on stock returns. Higher oil prices lead to higher production costs, thereby putting downward pressure on the value of assets. Sadorsky also emphasizes the importance (and asymmetry) of oil price volatility shocks in accounting for fluctuations in real stock returns.

Sawyer and Nandha (2006) consider the question of whether oil prices are a global factor in asset returns. They conclude that the effect of oil on stock prices is both disaggregated and heterogeneous. Indeed, Abdelaziz *et al.* (2008) examine the impact of oil prices on the stock markets of four Middle-East countries and find that oil prices have a strong and positive impact on these markets. Driesprong *et al.* (2008) find that higher oil prices result in lower stock returns in the following month. This negative relationship between oil prices and stock returns at a lag is significant both statistically and economically for many countries and for the world market index. The authors interpret their finding using the 'underreaction hypothesis': it takes time for investors to fully react to changes in the oil prices and thus stock markets do not fully reflect these changes until later.

Miller and Ratti (2009) report a negative long-run relationship between oil prices and stock markets for the 1970s and 1990s for several OECD countries, but virtually no relationship for the 1980s and 2000s. Agusman and Deriantino (2008) found that oil price changes did not have a significant impact on

industry stock returns in Indonesia, a member of OPEC. In a recent U.S. study, Gogineni (2009) examined the impact of rising oil prices on the stock returns of a wide variety of industries. The author found diverse and heterogeneous results across industries with some showing a quite favourable impact from rising oil prices, whereas other industries, such as airlines, had a negative impact. It therefore appears that the extant literature has not yet reached a consensus with respect to the empirical relationship between oil prices and broad financial indicators, and the research conducted in our paper provides a useful approach to refining previous results.

The application of wavelets in the field of finance is relatively new. However, given the versatility of wavelets, increased applications in finance can be expected. As Raihan *et al.* (2005) have noted, wavelets are particularly helpful in analysing non-stationary price series. Wavelet decomposition has also been applied to irregular time series to better analyse intraday trading volume (Manchaldore *et al.* 2010).

A recent study of commodity futures by Elder and Jin (2009) used wavelet analysis to ascertain the degree of fractional integration in various commodities. Applying the wavelet transform to different commodity time series allows the authors to obtain the sample variance of the time series from the wavelet detail coefficients at each scale. Wavelet transforms completely preserve all of the energy of a time series and the contribution of energy at each time scale, allowing an analysis of both high-frequency and low-frequency variation (Walker 1999). Scale-by-scale analysis also allows for a more detailed examination of the relationship between futures prices of the S&P 500 and futures prices of oil. An example of wavelet scale-by-scale analysis applied to the crude oil market is the paper by Jammazi and Aloui (2010), where the authors use multi-resolution analysis to highlight periods of extreme oil price volatility.

16.3 Data

The data for this study are the daily futures closing prices for oil and S&P 500 over the 28-year period from March 30, 1983, through March 4, 2011, as reported by Price-data.com. This continuous price data set is based upon the futures price of the near month contract, as this contract generally has the most open interest and thus the greatest liquidity. Prior to expiration of the futures contract, the position is rolled over into the next near month contract. Rolling over the contract prior to expiration allows Price-data.com to maintain a continuous time series of futures prices. In practice, the front month contract is rolled to the next month when its open interest is greater than the contract about to expire. The time period of study includes a wide variety of market conditions, such as the market crash of October 1987, the tech bubble of 2000, and the recent stock market volatility through the beginning of 2011. The time period of study also includes the period of extreme volatility in the oil market preceding the First Gulf War in spring 1991 as well as the very volatile year of 2008. In total, we have 7045 data points in our study. Figure 16.1 presents our data in levels (first two graphs), as well as log-differences (last two graphs). Panel (a) of Figure 16.1 depicts oil futures closing prices, while panel (b) presents S&P 500 futures prices. Panel (c) of Figure 16.1 shows returns on oil futures, and, finally, panel (d) plots returns on S&P 500 contracts. A review of the descriptive statistics in Table 16.1 shows a high degree of skewness and kurtosis.

Calculation of the traditional Pearson correlation coefficient is presented in Table 16.2 and corroborates the results of both wavelet analysis and Fourier analysis. The correlation matrix in Table 16.2 shows that, over our 28-year period of study, there is a pronounced positive correlation between crude oil futures and S&P 500 futures prices. In Table 16.3 we show the cross-correlations over 22 trading days (one calendar month of data). Again, we note the pronounced and persistently positive association between crude oil and the S&P 500. Finally, in Figure 16.2 we plot the cross-correlation function for ± 22 lags (number of trading days in a month) and show very strong statistical significance at all lags as all of the plotted values are far beyond the confidence limits for 99.9%.




FIGURE 16.1 (a, b) Daily futures closing prices and (c, d) daily futures returns.

TABLE 16.1 Descriptive Statistics: S&P 500 and Oil Futures Daily Prices and Per Cent Daily Returns

Variable	Oil Prices	Oil Returns	S&P 500 Prices	S&P 500 Returns
Mean	34.1679	0.0178	766.5020	0.0305
Median	25.0050	0.0444	770.0000	0.0621
Maximum	145.8600	13.3403	1576.2000	17.7493
Minimum	10.4200	-38.4071	148.8500	-33.7004
Std. dev.	23.8764	2.1757	443.4322	1.2654
Skewness	1.7548	-0.9477	0.1137	-2.4779
Kurtosis	5.7047	20.6091	1.4951	87.4236
Jarque-Bera	5762	92063	680	2099081
Probability	0.000	0.000	0.000	0.000
Observations	7044	7044	7044	7044

TABLE 16.2 Correlation Matrix

	Price _{Oil}	Price _{S&P}	Volume _{Oil}	Volume _{S&P}	Return _{Oil}	Return _{S&P}	Δ Volume _{Oil}	Δ Volume _{S&P}	Return _{Oil} ²	Return _{S&P} ²
Price _{Oil}	1	0.564***	0.740***	-0.330***	0.025**	-0.018	0.011	0.012	0.004	0.011
	[0]	[0.000]	[0.000]	[0.000]	[0.038]	[0.131]	[0.380]	[0.322]	[0.743]	[0.370]
Price _{S&P}	0.564***	1	0.562***	0.099***	0.022**	-0.002	0.003	0.016	0.010	-0.005
	[0.000]	[0]	[0.000]	[0.000]	[0.070]	[0.902]	[0.798]	[0.190]	[0.391]	[0.691]
Volume _{Oil}	0.740***	0.562***	1	-0.191***	0.008	-0.030**	0.009	0.009	0.059***	0.028**
	[0.000]	[0.000]	[0]	[0.000]	[0.490]	[0.012]	[0.434]	[0.446]	[0.000]	[0.020]
Volume _{S&P}	-0.330***	0.099***	-0.191***	1	-0.019	-0.064***	0.002	0.068***	0.026**	0.076***
	[0.000]	[0.000]	[0.000]	[0]	[0.120]	[0.000]	[0.842]	[0.000]	[0.030]	[0.000]
Return _{Oil}	0.025**	0.022**	0.008	-0.019	1	0.048***	0.008	-0.019	-0.210***	-0.022**
	[0.038]	[0.070]	[0.490]	[0.120]	[0]	[0.000]	[0.503]	[0.117]	[0.000]	[0.065]
Return _{S&P}	-0.018	-0.002	-0.030**	-0.064***	0.048***	1	0.003	-0.041***	0.029**	-0.262***
	[0.131]	[0.902]	[0.012]	[0.000]	[0.000]	[0]	[0.822]	[0.001]	[0.016]	[0.000]
Δ Volume _{Oil}	0.011	0.003	0.009	0.002	0.008	0.003	1	0.008	0.012	-0.003
	[0.380]	[0.798]	[0.434]	[0.842]	[0.503]	[0.822]	[0]	[0.480]	[0.300]	[0.823]
Δ Volume _{S&P}	0.012	0.016	0.009	0.068***	-0.019	-0.041***	0.008	1	0.007	0.006
	[0.322]	[0.190]	[0.446]	[0.000]	[0.117]	[0.001]	[0.480]	[0]	[0.582]	[0.639]
Return _{Oil} ²	0.004	0.010	0.059***	0.026**	-0.210***	0.029**	0.012	0.007	1	0.037***
	[0.743]	[0.391]	[0.000]	[0.030]	[0.000]	[0.016]	[0.300]	[0.582]	[0]	[0.002]
Return _{S&P} ²	0.011	-0.005	0.028**	0.076***	-0.022**	-0.262***	-0.003	0.006	0.037***	1
	[0.370]	[0.691]	[0.020]	[0.000]	[0.065]	[0.000]	[0.823]	[0.639]	[0.002]	[0]

P-Values in brackets. *, ** and *** denote 10%, 5% and 1% significance levels, respectively.

TABLE 16.3 Cross-Correlation Matrix

	Lag:	Lead:
<i>i</i>	Oil,S&P 500($-i$)	Oil,S&P 500($+i$)
0	0.563	0.563
1	0.564	0.562
2	0.564	0.561
3	0.564	0.560
4	0.565	0.559
5	0.565	0.559
6	0.565	0.558
7	0.566	0.557
8	0.566	0.556
9	0.566	0.555
10	0.566	0.554
11	0.567	0.553
12	0.567	0.552
13	0.567	0.551
14	0.567	0.551
15	0.567	0.550
16	0.568	0.549
17	0.568	0.548
18	0.568	0.547
19	0.568	0.546
20	0.568	0.546
21	0.569	0.545
22	0.569	0.544




FIGURE 16.2 Cross-correlations. The variance of the cross-correlation coefficient under the null hypothesis of zero correlation is approximately $1/n$, where n is the length of the series. Thus the blue boundary lines in the graph denote 99.9% confidence limits: $4(1/\sqrt{n}) = 0.047653$.

16.4 Methodology

16.4.1 Cross-Spectral Analysis

We conduct a cross-spectral analysis of the daily futures price of crude oil and the futures price of the S&P 500 contract between March 1983 and March 2011. We use frequency-domain techniques to determine the relative importance of cycles of different frequencies in accounting for the comovement among prices and returns on oil futures and the S&P 500. Further, we study how cospectra (which are the real component of the cross-spectra) change during two periods—before and after 1996 (i.e. the first and second halves of our sample). The goal of spectral analysis is to determine how important cycles of different frequencies are in accounting for volatility and comovement of the time series (Hamilton 1994). According to Granger (1966), one of the advantages of spectral methods is that they do not require specification of a model and so the results are not based on any rigid modelling assumptions.

To perform cross-spectral analysis, we first use the finite Fourier transform to decompose the data into a sum of sine and cosine waves of different amplitudes and wavelengths to obtain periodograms. According to the spectral representation theorem (Hamilton 1994), any covariance-stationary process x_t can be expressed as the finite Fourier transform decomposition of x_t :

$$x_t = \bar{x} + \sum_{k=1}^m [a_k \cos(\omega_k t) + b_k \sin(\omega_k t)], \quad (16.1)$$

where t is the time subscript ($t = 1, 2, \dots, n$), n is the number of observations in the time series, $\bar{x}$ is the mean value of x , m is the number of frequencies in the Fourier decomposition,^{*} a_k are the cosine coefficients, b_k are the sine coefficients, and ω_k are the Fourier frequencies ($\omega_k = 2\pi k/n$). Such an approach allows one to describe the value of x_t as a sum of periodic functions of different amplitudes and wavelengths. Each time series is thus decomposed into a number of orthogonal components associated with various frequencies.[†] We then calculate the amplitude cross-periodograms for each pair of time series as

$$J_k^{xy} = \frac{n}{2}(a_k^x a_k^y + b_k^x b_k^y) + i \frac{n}{2}(a_k^x b_k^y - b_k^x a_k^y), \quad (16.2)$$

where i represents the imaginary unit $\sqrt{-1}$. A cross-periodogram is a sample analogue of the population cross-spectrum and it shows the contribution of the k th harmonic to the total covariance between two data series.

The cross-periodogram is, admittedly, a volatile and inconsistent estimator of the cross-spectrum. In addition, it does not become more accurate with an increase in sample size. To overcome this deficiency, cospectral density estimates are produced by smoothing the real part of the cross-periodogram ordinates.[‡] The idea behind such non-parametric (or kernel) estimation is to use frequencies $\{\omega_k, \omega_{k\pm 1}, \omega_{k\pm 2}, \dots, \omega_{k\pm h}\}$ in estimating the cospectrum at ω_k . So the bandwidth parameter h and the relative weights (that must sum to unity) given to each frequency fully characterize the kernel. Because the kernel estimate is an average over a number of frequencies, and because estimates of the cospectrum at ω_k and ω_l are approximately independent for large n and $k \neq l$ (Hamilton 1994), kernel estimates are less volatile and provide more consistent estimates of the cross-spectrum than the cross-periodogram. We use a triangular weight function (with $h = 21$) in the moving average applied to the cross-periodogram to form smoothed cospectral density estimates.[§]

^{*} $m = n/2$ if n is even, and $m = (n-1)/2$ if n is odd.

[†] Since the value of $\cos(\omega t)$ repeats itself every $2\pi/\omega$ periods, a frequency ω corresponds to a period of $2\pi/\omega$.

[‡] There is an obvious trade-off: smoothing reduces the variance of the estimator but introduces a bias.

[§] The main conclusions of this paper are immune to using other relative weights (or kernels) applied to the periodogram ordinates to form the cospectral density estimates.

Finally, we compare the cospectra for all pairs of data series before and after 1996 (which is the midpoint of our sample). The cross-spectrum $s_{xy}(\omega)$ integrates to the unconditional covariance, and the quadrature spectrum $q_{xy}(\omega)$ (the imaginary part of the cross-spectrum) integrates to zero since $q_{xy}(-\omega) = -q_{xy}(\omega)$ (Hamilton 1994). Therefore, the area under the cospectrum (or the real component of the cross-spectrum) is equal to the covariance between x and y . Note that the cospectrum may be positive over some frequencies and negative over others.

16.4.2 Wavelets

Wavelets are small waves. By design, the distance from the peak of the wave to average sea-level is exactly offset by the distance from average sea-level to the trough of the wave. Mathematically, this admissibility condition is written as

$$\int_{-\infty}^{+\infty} \Psi(t) dt = 0. \quad (16.3)$$

An additional condition is that the energy of the series being studied must be preserved. This second condition is written as (Gençay *et al.* 2002)

$$\int_{-\infty}^{+\infty} |\Psi(t)|^2 dt = 1. \quad (16.4)$$

In the field of signal analysis, the French mathematician Joseph Fourier discovered that complex signals could be broken down into a sum of sine waves of differing frequencies. However, as a time series is split into an increasing number of frequencies, it becomes near impossible to tell at what point in time a sudden spike in a given frequency may have occurred.

In the development of wavelet analysis, the French seismologist Jean Morlet was particularly interested in maintaining the time point location of unusual breaks in a signal. His solution was that instead of splitting a signal into a number of different frequencies, he would keep the signal intact and change the length of the analysing wavelet (often by powers of 2, i.e. dyadic) (Burke-Hubbard 1998). Changing the length of the wavelet at each level of resolution gives rise to the phraseology: multiresolution analysis (MRA) (Mallat 1999). Wavelets of relatively short length are particularly good at capturing the high-frequency components of a time series, whereas longer length wavelets capture the low-frequency longer enduring components of a time series. As one proceeds from one level of resolution to the next, a graphical summary of the pronounced information effects is often presented as a wavelet power spectrum that is made up of the squared absolute value of the wavelet coefficients at each level of resolution.

The Morlet wavelet is given by

$$\Psi(t) = \frac{1}{\sqrt{2\pi}} e^{-i\omega t} e^{-t^2/2}, \quad (16.5)$$

where $i = \sqrt{-1}$, ω is the central frequency (most often set to $\omega = 6$ or $\omega = 2\pi$), and t is time. In order to implement the Morlet wavelet, one convolves a time series $x(t)$ with the wavelet. The continuous wavelet transform by length (translation) and scale (dilation) is then obtained (Gençay *et al.* 2002):

$$W(s, u) = x(t) \star \Psi_{u,s}(t), \quad (16.6)$$

where $\star$ indicates convolution, u is length, s is scale, and $\Psi_{u,s}(t) = (1/\sqrt{s})\Psi[(t-u)/s]$. We apply the wavelet transform to the time series for daily oil futures and S&P 500 futures. For our study, we are interested in the association both within and across the data series over a wide range of scales.

A wavelet transform may be either discrete or continuous. The discrete wavelet transform (DWT) gives the most compact representation of a signal, but because its wavelet spectrum in turn contains discrete blocks, it is difficult to examine the evolution of changing frequencies over time (Wang and Sassen 2008). Although the Morlet wavelet is a widely used analysing wavelet, other possibilities might include the ‘Paul’ or the ‘Derivative of Gaussian’ wavelet. The Derivative of Gaussian includes only real values and thus would not provide important information with respect to different data series being in or out of phase with each other. A more serious alternative to the Morlet wavelet would be the Paul wavelet. Based upon an important comparison paper by De Moortel *et al.* (2004), the Morlet wavelet has favourable characteristics with respect to capturing and tracking different frequency oscillations, while the Paul wavelet has an advantage with respect to time localization. Given that the time location of our daily data series is already known and that we are particularly concerned with tracking changes in frequency oscillations (especially large changes), we believe that the Morlet wavelet is more appropriate for our work. Addison (2005) also considered several different wavelet transforms and selected the Morlet wavelet given its versatility and widespread acceptance in many different disciplines.

16.4.3 Cross Wavelets

Following Torrence and Compo (1998), if time series X and Y have theoretical Fourier power spectra P_k^x and P_k^y , then the cross-wavelet distribution is given as

$$\frac{|W_n^x(s)W_n^y(s)|}{\sigma_x\sigma_y} \Rightarrow Z_v(p) \frac{\sqrt{P_k^x P_k^y}}{v},$$

where W_n^x and W_n^y are the n wavelet coefficients of X and Y , respectively, at scale s , σ_x and σ_y are the respective standard deviations, and P is the power spectra of the $k = 0, \dots, n/2$ frequency of X and Y .

For an autoregressive process of lag(1), α is the autocorrelation coefficient and

$$P_k = \frac{1 - \alpha^2}{1 + \alpha^2 - 2\alpha \cos[2\pi(k/N)]}.$$

For $v = 1$ (real wavelets), $Z(95\%) = 2.182$, and for $v = 2$ (complex wavelets), $Z(95\%) = 3.999$.

16.4.4 Cross Wavelet Phase Angles

A distinct advantage of working with a complex wavelet such as the Morlet wavelet transform is that you receive both real and imaginary results. While the real values are important for calculating cross correlations across different time scales, imaginary values provide important information with respect to whether or not two time series are in phase with each other and, if they are not, measurement of the circular angular difference between the two series can be used to ascertain which time series leads the other, and with wavelets we can determine this by both scale and time location.

From Zar (1999) and Grinsted *et al.* (2004), the mean of a circular set of angles for $(a_i, i = 1, \dots, n)$ is defined as

$$a_m = \arg(X, Y), \quad \text{with } X = \sum_{i=1}^n \cos(a_i) \text{ and } Y = \sum_{i=1}^n \sin(a_i),$$

and the circular standard deviation is

$$s = \sqrt{-2 \ln \frac{R}{n}},$$

where

$$R = \sqrt{X^2 + Y^2}$$

16.5 Empirical Results

Our paper analyses daily futures data from March 30, 1983, to March 4, 2011. We use frequency-domain techniques to determine the relative importance of cycles of different frequencies in accounting for the comovement among prices and volume on crude oil and the S&P 500 futures contracts. Further, we study if—and how—cospectra change from the first to the second half of the sample (after 1996) as oil price volatility had become more pronounced.

Wavelet transforms are adept at capturing short-term high-frequency fluctuations as well as pronounced long-term trends. Our analysis across time scales offers a comprehensive view of multiple interactions in the futures contracts of oil and the S&P 500 futures contracts. By working with a complex wavelet, such as the Morlet wavelet, we also analyse the degree to which the time series are in phase with each other.

16.5.1 Cospectral Densities

All of our time series are non-stationary in levels, but stationary in log-differences, as evidenced by unit root tests.* Therefore, legitimate cross-spectral results can be obtained only when the methodology is applied to log-differences (or returns). We apply the cross-spectral methodology summarized in the previous section to three pairs of time series: cospectral densities of (1) oil price returns and oil volume changes, (2) oil price returns and S&P 500 returns, and (3) oil volume changes and S&P 500 returns. [Figures 16.3](#) and [16.4](#) plot the cospectral densities of log-differences of our time series against frequency.

Figure 16.3 presents cospectral densities of oil price returns and volume (solid line), oil price returns and S&P 500 returns (dotted line), and changes in oil volume and S&P 500 returns (dashed line). It is easy to see that the correlation between log differences of oil prices and the stock market is close to zero at virtually all frequencies. This is an interesting result in light of the capital market and cash-flow theories summarized in [Sections 16.1](#) and [16.2](#). Correlation for the other two pairs of series—oil prices and volume and oil volume and S&P 500—is higher in absolute value, but positive and negative association is distributed fairly evenly across most frequencies. We observe a slightly stronger relationship between oil volume and S&P 500 over higher frequencies, namely for $\omega \geq 1.5$ (corresponding to the periodicity of four trading days or less).

Next, we are interested in analysing how the correlations among the three time series (oil price returns, oil volume changes and S&P 500 returns) changed over time and, more specifically, after the 1980s when oil price volatility became more pronounced. To this end, we compare the cospectral densities for two subsamples: from March 1983 to December 1996, and from January 1997 to March 2011. Panels (a) and (b) of Figure 16.4 report our results. The most important difference between the two subsamples is that covoatility along most of the frequency components for the oil price–oil volume pair and the oil volume–S&P pair is almost one order of magnitude higher during the second half of the sample. This is consistent with the higher overall volatility in the oil and stock markets after the 1980s.

* Augmented Dickey–Fuller and Phillips–Perron tests were performed for oil futures contract closing prices and S&P 500 futures contract closing prices. The hypothesis of a unit root in levels could not be rejected even at the 10% significance level; log-differencing the data, on the other hand, yields stationary time series at the 1% significance level.




FIGURE 16.3 Cospectral densities of oil price returns and oil volume changes, oil price and S&P 500 returns, oil volume changes and S&P 500 returns, 1983–2011.




FIGURE 16.4 Cospectral densities of oil price returns and oil volume changes, oil price and S&P 500 returns, oil volume changes and S&P 500 returns, 1983–1996 (a) and 1997–2011 (b).

Panel (b) suggests that positive covolatility between oil price and oil volume in the post-1996 subsample is concentrated in the relatively low frequencies of $\omega \leq 1.2$, which corresponds to cycles with a duration of more than five working days.

There are no other striking differences between the two subsamples. As with the entire sample, we observe both positive and negative covariances between log-differences of oil prices and S&P 500 contract returns. Both panels of Figure 16.4 suggest that the correlations between oil contract returns and S&P returns are virtually zero at all frequencies. Thus, the fact that we observed more negative changes in the post-1980s period relative to the 1980s did not affect the lack of correlation between log differences of oil prices and the stock returns.*

16.5.2 Wavelets

Figures 16.5 through 16.11 report the results of the wavelet analysis. Figure 16.5 depicts the continuous wavelet transform for the S&P 500. We note a pronounced effect between the scales at 64 and 128 over the observation period from 4500 to 5000. This indicates a lasting impact over three to six months (there are approximately 22 trading days per month). This impact occurred during calendar years 2000 and 2001 when the tech bubble rose to its highest peak and subsequently burst. The contour lines mark the 95% statistical significance, with the lighter shaded areas indicating the strongest impact.

Figure 16.6 plots the continuous wavelet transform for the oil futures prices. We note that the two most significant effects occur between scales 128 and 256 (6 to 12 months) over the observation periods 2000 and 6250. These periods correspond to the first time that the U.S. was at war with Iraq and is indicative of the disruption to the oil markets at that time, as well as the huge speculative oil bubble




FIGURE 16.5 Continuous wavelet transform—S&P 500 futures closing prices.

* As with the analysis of the entire sample, we find that, quite predictably, oil futures volume is more closely associated with oil prices than with S&P 500 returns.




FIGURE 16.6 Continuous wavelet transform—Oil futures closing prices.




FIGURE 16.7 Three-dimensional view of the wavelet power spectrum—Oil futures.




FIGURE 16.8 Three-dimensional view of the wavelet power spectrum—S&P 500 futures.




FIGURE 16.9 Cross-correlations: Oil prices and oil volume.




FIGURE 16.10 Cross-correlations: Oil price and S&P 500 close.




FIGURE 16.11 Cross-correlations: S&P 500 and S&P 500 volume.

buildup and collapse in 2008. Again, the contour lines mark the 95% statistical significance, with the lighter shaded areas indicating the strongest impact.

Figure 16.7 plots a three-dimensional view of the wavelet power spectrum (squared wavelet detail coefficients). The wavelet power spectrum illustrates the pronounced effect at observation period 2000 corresponding to the first U.S. war with Iraq and its impact on the oil market, and an even more pronounced impact at approximately observation number 6250 corresponding to the dramatic increases in oil prices in summer 2008, followed by the dramatic decline in what many considered to be a huge speculative bubble. In Figure 16.8, the wavelet power spectrum for S&P 500 shows the most pronounced effect occurring at approximately observation period 4500, which corresponds to the tech bubble and its subsequent collapse at the turn of the century. Also notice the strongest scale effect at 1024 (trading days) corresponding to the 4-year presidential impact on stock prices.

In Figure 16.9, the dark bold contour lines mark the 95% significance threshold level where we see the strongest wavelet correlation between the futures price of oil and the futures trading volume of oil. The cross-wavelet correlation is strongest at 128 trading days (6 months) and at 256 trading days (1 year).

Figure 16.9 also shows the statistically significant phase angles plotted by directional arrows. An arrow pointing to the right indicates that both time series are exactly in phase with each other, whereas an arrow pointing to the left indicates that the two time series are completely out of phase with each other. An arrow pointing straight up indicates that the first data series (oil futures price) leads the second data series (oil futures volume) and an arrow pointing straight down indicates the second series leads the first. As we can see from Figure 16.9, the arrows predominantly point up, indicating that oil prices lead oil volume, or they point to the right, indicating that the two data series are in phase with each other. Figure 16.9 suggests that oil prices lead oil volume for about 256 days (one year) in the early 2000s.

In Figure 16.10, we observe the wavelet cross-correlations between oil prices and S&P 500 prices. The strongest field of correlation is found in the earlier part of our study (approximately the first time the U.S. went to war in Iraq). The directional phase angle arrows predominately point up and to the left in the earlier part of our study, indicating a mix between oil prices leading S&P 500 futures prices and S&P 500 futures prices being out of phase with oil futures prices. Towards the end of the study, the phase arrows point to the right, indicating that the two series move together. Note that the movement together occurs over a wide range of scales.

Finally, in Figure 16.11, the predominate cross-correlations occur over a 128 trading day period (6 months). Note that the statistically significant directional phase angle arrows mostly point down, indicating that increased trading volume leads to higher S&P futures prices by approximately 6 months. Arrows that are pointing to the left in this figure indicate that volume and prices are out of phase.

16.6 Conclusions

Our results indicate that, over the 28-year period of study, there is a persistent positive association between crude oil futures prices and S&P 500 futures prices. This finding is observed over a large number of time scales and over a very wide range of prices. With respect to returns, cross-spectral analysis indicates that there is virtually no contemporaneous relationship between oil futures returns and S&P 500 returns at most frequencies. The positive correspondence between oil futures prices and S&P 500 futures prices is contrary to the purported negative relationship between oil prices and the U.S. stock market reported by Sauter and Awerbuch (2003) and earlier by Jones and Kaul (1996).

Our seemingly counterintuitive findings may be due to the likely presence of nonlinearities in the relationship between the oil market and the macroeconomy suggested by the extant literature (e.g. Lee *et al.* (1995) and Hamilton (1996)). Tang and Xiong (2010), who find that commodity prices indeed became increasingly more correlated with equity markets after 2000, point to the rise of investments in commodity futures in the mid-2000s, which suggests that commodity prices can be affected by

financial factors in addition to changes in supply and demand. Another possible explanation lies in the bi-directional causality between oil futures returns and the S&P 500 returns emphasized by Ciner (2001). The cross-spectral and wavelet analyses employed in our paper allow for nonlinear feedback effects between oil prices and financial markets.

We also analyse the relationship between futures prices and trading volume. The directional phase angle plots reported in the empirical section suggest that oil prices lead oil volume and S&P 500 trading volume leads S&P 500 futures prices. In particular, oil prices lead oil volume for about 256 days (one year) in the early 2000s. Generally speaking, our results can be viewed as reflective of the investors' bounded rationality and the gradual information diffusion hypothesis (Hong and Stein 1999). Our empirical analysis indicates that investors and policy makers should pay close attention to trading volume when analysing stock market movements.

Conventional wisdom and past research (see, for example, Sauter and Awerbuch (2003) as well as works cited therein) suggest that changes in oil prices should be negatively related to macroeconomic and financial indicators. Although it is tempting to assume that higher oil prices bode ill for companies' stock valuations, our paper shows that this assumption is spurious. Contrary to the conventional wisdom, our findings point to the positive association between oil and S&P 500 futures for a large number of time scales and across a wide range of Fourier frequencies. This finding of a positive association between oil prices and stock markets is logical from an economic standpoint. When the economy is recovering from a recession and economic conditions start to improve, stock prices generally go up. During an expansion, wealth is positively affected through higher net worth of the companies and the corresponding higher value of personal assets. Consumers' incomes also become higher due to better and more abundant employment opportunities, as well as due to higher wages during an expansion. Higher demand for consumer and industrial goods due to increased public wealth and incomes, coupled with better balance sheets of companies, lead to a greater demand for energy, which in turn drives up prices of oil. Thus, higher economic activity is associated with higher oil prices. Conversely, when the economy begins to falter, stock market valuations of firms become lower. These negative movements in the stock markets decrease public wealth and companies' net worth, thereby further depressing economic activity. Lower levels of economic activity, consumers' wealth and companies' net worth, as well as the ensuing weaker demand for energy, put downward pressure on oil prices and volume. The main findings of our paper are, therefore, consistent with economic theory.

This paper finds no evidence of disruptive effects of oil price movements on the financial markets. Although there may be many reasons to switch from fossil fuels such as oil to alternative forms of renewable energy (e.g. balance of payments problem, rising concentrations of CO₂, dependence on unstable regimes), our results indicate that an unfavourable impact on U.S. equity markets from rising oil prices is not supported by empirical research.

The results reported in this paper have important practical implications for investors: it is possible to improve the accuracy of the forecast of future returns by studying oil futures prices and futures volumes. Unlike most of the literature that examines the relationship between oil prices and the stock market, we use daily data on oil futures and S&P 500 futures, which allows us to conduct a more detailed analysis of the said relationship. Careful empirical investigations of oil-stock market relationships can have important implications for portfolio selection (e.g. Arouri and Nguyen (2010)). Finally, the frequency-domain techniques used in this paper help circumvent the issue of possible endogeneity of oil prices with respect to the U.S. economy that often plagues the literature's attempts to discern the impact of oil price changes on the U.S. stock market (Kilian 2009, Kilian and Park 2009).

Acknowledgements

We would like to thank two anonymous referees for invaluable comments. A grant for futures data is gratefully acknowledged; information is available at <http://www.price-data.com/>.

Software Acknowledgements

1. IDL (Interactive Data Language), <http://www.ittvis.com/ProductServices/IDL.aspx>
2. Grinsted, A., Moore, J.C. and Jevrejeva, S., <http://www.pol.ac.uk/home/research/waveletcoherence/>
3. Torrence, C. and Compo, G., <http://atoc.colorado.edu/research/wavelets/>
4. S W A N (Software for Waveform Analysis), CNRS/LPCE, Orleans, France, Lagoutte, D., Brochot, J.Y. and Latremoliere, P., <ftp://lpce.cnrs-orleans.fr/projects/swan/pub/swan-2.42/doc/swan-at.pdf>

References

- Abdelaziz, M., Chortareas, G. and Cipollini, A., Stock prices, exchange rates, and oil: Evidences from Middle East oil-exporting countries. In *Proceedings of the Middle East Economic Association*, New Orleans, LA, January 4–7, 2008.
- Addison, P.S., Wavelet transforms and the ECG: A review. *Physiol. Measure.*, 2005, **26**(5), R155–R199.
- Agusman, A. and Deriantino, E., Oil price and industry stock returns: Evidence from Indonesia. In *21st Australasian Finance and Banking Conference*, August 25, 2008.
- Arouri, M.E.H. and Nguyen, D.K., Oil prices, stock markets and portfolio investment: Evidence from sector analysis in Europe over the last decade. *Energy Policy*, 2010, **38**(8), 4528–4539.
- Awerbuch, S., The surprising role of risk and discount rates in utility integrated-resource planning. *Electric. J.*, 1993, **6**(3), 20–33.
- Awerbuch, S. and Deehan, W., Do consumers discount the future correctly? A market-based valuation of fuel switching. *Energy Policy*, 1995, **23**(1), 57–69.
- Bolinger, M., Wiser, R. and Golove, W., Quantifying the value that wind power provides as a hedge against volatile natural gas prices. Lawrence Berkeley National Laboratory Working Paper, 2002.
- Burke-Hubbard, B., *The World According to Wavelets: The Story of a Mathematical Technique in the Making*, 1998 (AK Peters Ltd).
- Campbell, J.Y., A variance decomposition for stock returns. *Econ. J.*, 1991, **101**, 157–179.
- Ciner, C., Energy shocks and financial markets: Nonlinear linkages. *Stud. Nonlin. Dynam. Econometr.*, 2001, **5**(3), 203–212.
- De Moortel, I., Munday, S.A. and Hood, A.W., Wavelet analysis: The effect of varying basic wavelet parameters. *Solar Phys.*, 2004, **222**(2), 203–228.
- Driesprong, G., Jacobsen, B. and Maat, B., Striking oil: Another puzzle? *J. Financ. Econ.*, 2008, **89**(2), 307–327.
- Elder, J. and Jin, H.J., Fractional intergration in commodity futures returns. *Financ. Rev.*, 2009, **44**(4), 583–602.
- Elder, J. and Serletis, A., Oil price uncertainty. *J. Money, Credit Bank.*, 2010, **42**(6), 1137–1159.
- Ferderer, J.P., Oil price volatility and the macroeconomy. *J. Macroecon.*, 1996, **18**(1), 1–26.
- Gençay, R., Selcuk, F. and Whitcher, B., *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*, 2002 (Academic Press: New York).
- Gogineni, S., Oil and the stock market: An industry level analysis. Working Paper Series, University of Oklahoma, Division of Finance, 2009.
- Granger, C.W.J., The typical spectral shape of an economic variable. *Econometrica*, 1966, **34**(1), 150–161.
- Grinsted, A., Jevrejeva, S. and Moore, J., Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlin. Process. Geophys.*, 2004, **11**, 561–566.
- Hamilton, J.D., Oil and the macroeconomy since World War II. *J. Polit. Econ.*, 1983, **21**(2), 228–248.
- Hamilton, J.D., *Time Series Analysis*, 1994 (Princeton University Press: Princeton).
- Hamilton, J.D., This is what happened to the oil price–macroeconomy relationship. *J. Monet. Econ.*, 1996, **38**(2), 215–220.
- Hamilton, J.D., What is an oil shock? NBER Working Paper 7755, 2000.

- Hong, H. and Stein, J.C., A unified theory of underreaction, momentum trading and overreaction in asset markets. *J. Finance*, 1999, **54**(6), 2143–2184.
- Hooker, M.A., What happened to the oil price–macroeconomy relationship? *J. Monet. Econ.*, 1996, **38**(2), 195–213.
- Huang, R.D., Masulis, R.W. and Stoll, N.R., Energy shocks and financial markets. *J. Fut. Mkts*, 1996, **16**(1), 1–27.
- Jammazi, R. and Aloui, C., Wavelet decomposition and regime shifts: Assessing the effects of crude oil shocks on stock market returns. *Energy Policy*, 2010, **38**(3), 1415–1435.
- Jones, C.M. and Kaul, G., Oil and the stock markets. *J. Finance*, 1996, **51**(2), 463–91.
- Kilian, L., Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *Am. Econ. Rev.*, 2009, **99**(3), 1053–1069.
- Kilian, L. and Park, C., The impact of oil price shocks on the U.S. stock market. *Int. Econ. Rev.*, 2009, **50**(4), 1267–1287.
- Lee, C.-C. and Zeng, J.-H., The impact of oil price shocks on stock market activities: Asymmetric effect with quantile regression. *Math. Comput. Simul.*, 2011, **81**(9), 1910–1920.
- Lee, K., Ni, S. and Ratti, R.A., Oil shocks and the macroeconomy: The role of price variability. *Energy J.*, 1995, **16**(4), 39–56.
- Mallat, S., *A Wavelet Tour of Signal Processing*, 2nd ed., 1999 (Academic Press: New York).
- Manchaldore, J., Palit, I. and Soloviev, O., Wavelet decomposition for intra-day volume dynamics. *Quantit. Finance*, 2010, **10**(8), 917–930.
- Miller, J.I. and Ratti, R.A., Crude oil and stock markets: Stability, instability, and bubbles. *Energy Econ.*, 2009, **31**(4), 559–568.
- Mork, K.A., Oil and the macroeconomy, when prices go up and down: An extension of Hamilton's results. *J. Polit. Econ.*, 1989, **97**(3), 740–744.
- Mork, K.A., Olsen, O. and Mysen, H.T., Macroeconomic responses to oil price increases and decreases in seven OECD countries. *Energy J.*, 1994, **15**(4), 19–35.
- Raihan, S.Md., Wen, Y. and Zeng, B., Wavelet: A new tool for business cycle analysis. Federal Reserve Bank of St. Louis Working Paper 2005-050, 2005.
- Reboredo, J.C., Nonlinear effects of oil shocks on stock returns: A Markov-switching approach. *Appl. Econ.*, 2010, **42**(29), 3735–3744.
- Rotemberg, J.J. and Woodford, M., Imperfect competition and the effects of energy price increases on economic activity. *J. Money, Credit, Bank.*, 1996, **28**(4), 550–577.
- Sadorsky, P., Oil price shocks and stock market activity. *Energy Econ.*, 1999, **21**, 449–469.
- Sauter, R. and Awerbuch, S., Oil price volatility and economic activity: A survey and literature review. IEA Working Paper, 2003.
- Sawyer, K.R. and Nandha, M., How oil moves stock prices. Working Paper Series, University of Melbourne, Department of Finance, 2006.
- Tang, K. and Xiong, W., Investment and nancialization of commodities. NBER Working Paper No. 16385, 2010.
- Torrence, C. and Compo, G.P., A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.*, 1998, **79**, 61–78.
- Vo, M., Oil and stock market volatility: A multivariate stochastic volatility perspective. *Energy Econ.*, 2011, **33**(3), 956–965.
- Walker, J.S., *A Primer on Wavelets and their Scientific Applications*, 1999 (CRC Press: Boca Raton).
- Wang, L. and Sassen, K., Wavelet analysis of cirrus multiscale structures from lidar backscattering: A cirrus uncinus complex case study. *J. Appl. Meteorol. Climatol.*, 2008, **47**(10), 2645–2658.
- Yang, C.W., Hwang, M.J. and Huang, B.N., An analysis of factors affecting price volatility of the US oil market. *Energy Econ.*, 2002, **24**, 107–119.
- Zar, J.H., *Biostatistical Analysis*, 1999 (Prentice Hall: Engelwood Cliffs, NJ).

Sectoral Stock Return Sensitivity to Oil Price Changes: A Double- Threshold FIGARCH Model

Elyas Elyasiani

Iqbal Mansur

Babatunde Odusami

17.1	Introduction	360
17.2	Data and Methodology	362
	Data Description • Methodology and Models • Testable Hypotheses	
17.3	Empirical Findings	369
	Tests of Model Specification • Tests of Asymmetry • Sector Analysis • The FIGARCH Volatility Equation	
17.4	Conclusions.....	383
	Acknowledgements.....	384
	References.....	384
	Appendix A: Threshold and Delay Parameter Values	386

We investigate the association between the stock return distributions of 10 major U.S. sectors and oil returns within a double-threshold FIGARCH model. This model nests GARCH, IGARCH and Fama–French specifications as its special cases and allows a test of their validity. This model also has the advantage of capturing not only the short-run dynamics (as in the standard GARCH model), but also the long-run persistence pattern of oil shock effects that may decay at a slower hyperbolic pace. We find that: (i) data reject the more restricted GARCH, IGARCH and Fama–French models in favour of FIGARCH, (ii) oil return is a significant determinant of every sector's return and/or return volatility, (iii) oil effects are asymmetric for oil returns above and below the thresholds, (iv) asymmetry is stronger when oil return volatility is greater, (v) volatilities of sectoral returns exhibit threshold-based regime shifts and (vi) oil shock effects follow a hyperbolic, rather than an exponential decay pattern, establishing long-term persistence of shocks. Policy implications are drawn.

Keywords: Oil price; Sector returns; Threshold; FIGARCH

JEL Classification: G1, G12

17.1 Introduction

Despite the substantial adverse effects of oil price shocks on the U.S. economy since the OPEC oil embargo of 1973, research on the association between oil price changes and the U.S. stock market has been relatively limited.* In particular, no study has been conducted on the threshold effect of the oil price changes on the equity return distribution at the sector level. Identification of threshold values for oil price effects is important because basic models overlooking this non-linearity distort the effects of oil shocks on both the stock market and the real economy. Moreover, the findings in the existing literature are conflicting in terms of significance, magnitude and the direction of the oil shock effects. The pioneering work of Chen *et al.* (1986) suggests that the risk associated with oil price changes was not priced in the U.S. stock market even during the 1968–1977 period when OPEC played a critical role in setting oil prices. However, Jones and Kaul (1996) demonstrate that changes in real oil returns did have a detrimental effect on output and real stock returns in the U.S., Canada, Japan and the U.K. Similarly, Huang *et al.* (1996) find that futures returns on heating and crude oil did lead individual oil company and petroleum industry stock returns during the 1980s, but failed to exert an impact on other industry stock indices or the S&P 500 index. Oil futures volatility was also found to lead the petroleum stock index volatility, but not other industry or market indices.

Sadorsky (1999) employs a Generalized Autoregressive Conditionally Heteroskedastic (GARCH) model to investigate the interaction of oil price shocks and the stock market. He also uses the vector auto regression (VAR) technique to examine the interaction between return on oil and industrial production, interest rate and real stock returns. His results provide evidence of unidirectional and asymmetric shock transmission from oil to stocks. Moreover, positive oil price shocks and oil price volatility shocks explain a larger portion of the forecast error variance in real stock returns than do their negative counterparts, at least for some sub-periods.

Faff and Brailsford (1999) investigate the sensitivity of the Australian industry equity returns to oil price changes over the period July 1983 to March 1996. Their results suggest that oil is an important factor in describing the return generating process of Australian industries. However, the direction of the effect is industry-specific; oil, gas and diversified resources industries show positive sensitivity, while paper and packaging, and transport industries demonstrate negative sensitivity to oil shocks.† Building on Sadorsky (1999), Huang *et al.* (2005) use a multivariate threshold model to investigate the impact of real oil return and its volatility on industrial production, interest rate and real stock returns for the U.S., Canada and Japan. Using monthly data from 1970 to 2002, they uncover the existence of differential thresholds for oil return and oil return volatility in their impacts on each of these three economies. The threshold values are found to be dependent on the extent to which an economy is reliant on imported oil. Their findings suggest that when oil return, or its volatility, are below their corresponding threshold levels, they have limited impacts on the macroeconomic variables considered, while if they are above the

* The impact of oil price shocks on the U.S. and the global economy is considerable. It is estimated that the oil price shocks of 2004–2008 cost the U.S. economy nearly \$1.9 trillion (<http://www.fueleconomy.gov/FEG/oildep.shtml>). Similarly, the International Energy Agency (IEA) estimates that for every \$1 increase in oil prices, the world GDP is adversely affected by \$25 billion (<http://www.iea.org>). Along these lines, Greene and Tishchishyna (2000) report that from 1970 to 2000, oil price movements resulted in a \$7 trillion loss in terms of costs to the U.S. Yang *et al.* (2002) also argue that higher oil prices engender recessions and unemployment in oil-consuming nations, as they curtail economic activity (Awerbuch and Sauter 2005). These findings have major implications for the formulation of the monetary, fiscal and energy policies of oil exporting/importing countries. Moreover, dependence on imported oil has major implications in terms of political leverage and political subordination, examples of which were clearly manifested in the aftermath of the 1973 oil embargo and the 1979 sharp increase in oil prices.

† The magnitude of the distortion due to model misspecification can be considerable (see footnote Section 17.2.3).

threshold levels, changes in oil returns demonstrate greater explanatory power than those in oil return volatility.

Given the vital importance of oil to the U.S. economy, the relationship between sectoral stock returns and oil futures returns, and sensitivity of the stock return mean and volatility to oil shocks require further scrutiny. The main goal of this study is to provide a detailed analysis of the association between the equity return distributions of different sectors in the U.S. economy and changes in the oil returns. To this end, we take the following steps: (i) we model the sectoral stock return effects of oil return changes within a generalized Double-Threshold Fractionally Integrated GARCH (DT-FIGARCH) framework, (ii) we determine whether there is a threshold level of oil returns for each of the sectors considered, beyond which the oil return effects on equity return levels change in character and magnitude; (iii) we investigate whether a threshold effect is evident in the conditional volatility of sector returns; and (iv) we examine the effect of oil shocks on the return of each sector and its volatility within this generalized framework.

The contributions of this study are as follows: (i) this is the first study to apply a double-threshold FIGARCH methodology in order to allow for long memory and to assess the threshold effects of oil movement on both the first and the second moments of the return distribution for a large number of sectors; (ii) we test for the existence of, and identify the values of the threshold oil returns for different sectors and establish the asymmetry of the responses of sector returns, and their volatilities, to oil return changes when the latter is above/below the threshold; (iii) we provide strong evidence of statistically significant long-memory processes in the conditional variances of sector returns and (iv) we establish that the longevity of shock memory in the conditional volatility of returns is sector-specific and responsive to threshold regime shifts.

Our findings have several policy implications. The primary advantage of modelling conditional variance using a FIGARCH model is that it not only captures the short-run dynamics (as in the standard GARCH model), but also the long-run persistence that decays at a slower hyperbolic pace. Therefore, this long-memory model is capable of carrying information over a year or longer and delineating its pattern of effect. The implication of the long memory is that policy makers should not solely consider the effects of the latest turmoil or extreme events in the formulation of overall or industry-specific policies; they should also account for the impacts of the events in the distant past in this regard. Additionally, if volatility decay exhibits a mean-reverting behaviour, policy makers may be able to take appropriate actions to smooth the effect of a shock, shorten its duration and accelerate the mean-reversion process. Examples of such policy actions include the release of oil from the U.S. Strategic Petroleum Reserve (SPR) in the aftermath of Hurricane Katrina, the 2008 suspension of oil purchases for SPR in order to ease the upward pressure on oil prices, tax credits used to encourage investment in renewable energy and tax surcharges introduced to reduce oil consumption. Correct understanding and appropriate modelling of stock return volatility is also important because it can aid investors in their quest for optimal portfolio selection and risk management.

Analysis of stock return behaviour at the sector level is appealing to stock investors because while globalization has reduced the benefits of international risk diversification, sectors across national markets continue to demonstrate independent movements and, hence, sector-based investment portfolios can be helpful in controlling risk. In this context, Moskowitz and Grinblatt (1999) find that investment strategies based on 'industry-momentum' produce holding-period returns that are substantially larger than 'individual stock' momentum-based strategies. Moreover, findings about the sector-specific risk and return effects of oil price changes have important implications for planning decisions by producers in these sectors and government assessment of cost/benefits of regulatory policies designed to promote the use of alternative energy sources. Similarly, understanding sector-based volatility is important for determining the sectoral cost of capital, determination of sectoral shock persistence and asset allocation decisions, because volatility changes are likely to alter these decisions and outcomes.

Extant studies mostly focus on the first moment of the return distribution, overlooking the second moment or the volatility equation. Given that financial time series are known to exhibit GARCH properties, the assumption of a fixed-variance error term in the mean equation results in erroneous inferences. Even studies that do account for GARCH properties generally adopt a single-regime scenario, overlooking the presence of non-linearities such as threshold and asymmetry effects. Given that the degree of volatility persistence is likely to be sensitive to whether the return on oil (R_{oil}) takes values above or below the threshold, these thresholds need to be recognized, identified and accounted for. Sadorsky (1999) arbitrarily chooses the zero change level as the threshold for all macro-variables he studies. In practice, however, idiosyncratic characteristics of economic sectors are likely to result in non-zero and dissimilar threshold values for different sectors. We remedy this shortcoming by adopting a Tsay (1998) type methodology to estimate sector-specific threshold values for each period studied. We specify volatility as a fractionally integrated (FIGARCH (p, δ, q)) model, developed by Baillie *et al.* (1996). This model adequately captures the decay mechanism of the autocorrelation function of the conditional volatility process, allowing one to properly analyse the long-term persistence of shocks to the volatility process. The remainder of the paper is organized as follows. Section 17.2 provides the data definitions and methodology. Empirical results are presented in Section 17.3 while the conclusions are given in Section 17.4.

17.2 Data and Methodology

17.2.1 Data description

Ten sectors, constructed based on two-digit SIC codes, are analysed in this study and roughly sorted into four major types:

- a. Oil substitutes; Coal Mining (SIC 12), Electric–Gas Services (SIC 49);
- b. Oil-related; Oil and Gas Extraction (SIC 13), Petroleum Refining (SIC 29);
- c. Oil-consuming; Building Construction (SIC15), Chemicals & Allied Products (SIC 28), Rubber & Plastic (SIC30), Transport Equipment (SIC 37) and Air Transportation (SIC 45);
- d. Financial; Depository Institutions (SIC 60).

Sector return data are collected from the Centre for Research in Security Prices (CRSP) database. The Fama–French factors used in the model include the market risk premium (RM), the size factor (SMB) and the value factor (HML) and are obtained from the Kenneth R. French website (www.mba.tuck.dartmouth.edu). SMB is the return on a portfolio long on small firms and short on large firms (small minus large firm returns). HML is the return on a portfolio long on stocks with high book-to-market ratio and short on stocks with low book-to-market ratio (high minus low market-to-book firm returns).

The focus of this study is the *ex ante* effect of oil price movements on industry returns. Hence, returns on crude oil futures are favoured over other oil time series as an indicator of oil return changes. Moreover, Sadorsky (2001) shows that spot prices of oil are more susceptible to transitory random noise than oil futures prices.* Following Guo and Kliesen (2005), the return on one-month crude oil futures traded on the New York Mercantile Exchange (NYMEX) is used as the oil return variable. The daily return on oil (R_{oil}) is calculated as $\log(p_t/p_{t-1})$ where p is the one-month crude oil futures price over the period of April 4, 1983 to December 29, 2006.

We distinguish between periods of tranquil and highly volatile oil market conditions because market reactions to shocks and, hence, threshold levels are likely to differ under these circumstances. The oil

* A host of oil price series are used in the literature: refiners' acquisition cost for crude oil (Mork 1989), PPI for oil (Jones and Kaul 1996, Sadorsky 1999), net oil price increase dummy (Lee and Ni 2002), crude petroleum price (Davis and Haltiwanger 2001, Hamilton 2003), PPI of oil relative to inflation (Hooker 2002), average of world spot crude oil prices deflated by county-specific PPI (Huang *et al.* 2005) and log difference of multiple oil series such as Arab Light, West Texas, Dubai, Brent, Brent Futures and NYMEX Futures (Driesprong *et al.* 2008). Due to space considerations a detailed discussion of this issue is not presented here. See Elyasiani *et al.* (2011) for more on the preference for oil futures returns over oil spot returns.

market volatility as measured by the standard deviation (SD) of oil prices was relatively high in the 1980s, low in the 1990s and very high in the 2000s. Specifically, the 1983–1991 sub-period exhibits a moderate level of volatility (SD of 6.31%) followed by relative calm in the 1991–1998 sub-period (SD of 2.71%) and extreme volatility in the 1998–2006 sub-period (SD of 15.49%). The corresponding oil price ranges (averages) for these three periods were \$10.42–40.42 (22.55), \$10.72–26.62 (19.15) and \$17.45–77.03 (39.81), respectively.*

The oil price data show that, since 1984, the oil market has seen two distinct structural shifts, one in 1990 and the other in 1998 (Figure 17.1). We employ the Zivot and Andrews (1992) (ZA) test to determine the exact location of the structural changes in this market. Using the full sample period, the ZA-test procedure identifies January 24, 1991 as the first break point. The test is repeated for the remaining observations and the second break point is found to be on December 10, 1998.[†] Estimation is carried out for these three distinct sub-periods separately. However, to save space, only the results of the latter two periods (tranquil and highly volatile, respectively) are presented. This allows us to contrast the threshold effects of oil return under these dissimilar environments. We will refer to the 1991–1998 and 1998–2006 as the first (low volatility) and second (high volatility) periods, respectively. The 1983–1991 period results are rather similar in scope to those obtained for the 1998–2006 period (available upon request).




FIGURE 17.1 NYMEX One Month Oil Futures Prices—April 4, 1983 to December 29, 2006.

* We report a similar set of statistics for oil futures return in Table 17.1.

[†] The advantage of the ZA procedure is that, unlike Perron (1989), no *a-priori* knowledge of the break points in the trend function is necessary. Instead, a data-dependent algorithm is used to determine the structural breaks. Thus, the algorithm produces unconditional unit-root test results. The ZA-test statistics for unit-root tests are -28.72 and -24.74 for the first and second structural break points, respectively. The critical value at the 1% level is 5.57 .

^{*} An unparalleled level of economic and political activity affected oil prices both in January 1991 and December 1998. A major factor that contributed to the global uncertainty of oil price movements was the first Gulf War (U.S. and allied countries' air strikes on Iraq beginning on January 16, Iraqi SCUD missile attacks on Israel on January 18 and 22, and the World's largest oil spill caused by the Iraqi forces in Kuwait on January 23). Following the Gulf War, crude oil prices entered a period of steady decline until 1994. The price cycle then turned up. The U.S. economy was strong and the Asian Pacific region was booming. Asian consumption accounted for all but 300,000 barrels per day of that gain and contributed to a price recovery that extended into 1997 (www.wtrg.com). The rapid growth in Asian economies came to a halt in 1998, leading to a lower consumption of oil. Higher OPEC production, coupled with lower consumption led to a sharp decline in oil prices, which continued through December 1998, recovered in early 1999 and continued to rise throughout 2000 (www.wtrg.com).

Table 17.1 presents the summary statistics for the sample period (1/25/1991 to 12/29/06).^{*} Panel A (panel B), reports the results for the 1991–1998 (1998–2006) sub-period. In the low volatility sub-period, the Depository Institutions sector had the highest return while Petroleum Refinery (oil-related sector) had the lowest return. In the same sub-period, the Coal Mining sector had the highest unconditional risk (SD 1.83%) of daily equity return, while Electric–Gas services had the lowest (SD 0.46%). For the high volatility sub-period (1998–2006), the data show that the Oil and Gas Extraction sector (an oil supplier industry) had the highest average daily return while the Air Transportation and Transport Equipment sectors (an oil user industry) had the lowest average daily return. The highest level of unconditional risk (SD 1.63%) is seen for the Building sector while the lowest (SD 0.52%) is for the Depository Institutions sector. On a risk-per-unit of return basis (coefficient of variation), the Petroleum Refinery sector had the highest risk in the low volatility sub-period, while Air Transportation had the highest in the high volatility sub-period. Depository Institutions had the lowest risk-per-unit of return in both sub-periods.

During the low volatility period, the skewness coefficients of daily returns for the Coal, Building, Petroleum Refinery, Plastic and Rubber and Air Transportation sectors were positive while those of Oil–Gas Extraction, Transport Equipment, Electric–Gas Services and Depository Institutions sectors were negative. In the high volatility sub-period, skewness is negative for all sector returns, except Coal and Building Construction. The skewness and kurtosis values and the Jarque–Bera (1987) (J–B) statistics are highly significant for all sectors in both sub-periods, indicating that the unconditional distributions of returns are non-normal. The Ljung–Box (1978) test statistics for the return series also reject the null hypothesis of no autocorrelation in both sub-periods, except for Coal Mining, Petroleum Refinery and Building sectors in the high volatility sub-period. Evidence of non-normality and the presence of serial correlation in the return distribution strongly support the application of a GARCH-type methodology.^{†‡}

17.2.2 Methodology and models

Our threshold framework of the oil effect on sector returns combines the FIGARCH (1, δ , 1) volatility specification of Baillie *et al.* (1996) with a four-factor generalization of Tsay (1998) to form a double-threshold FIGARCH (1, δ , 1) model (DT-FIGARCH):

$$\begin{aligned}
 R_{it} &= \sum_{j=1}^{J=2} I_t^j (b_0^{(j)} + b_1^{(j)} RM_t + b_2^{(j)} SMB_t + b_3^{(j)} HML_t \\
 &\quad + b_4^{(j)} R_{oil,t-d} + \epsilon_t), \epsilon_t \sim N(0, \sigma_t) | \Phi_t, \\
 \sigma_t^2 &= \sum_{j=1}^{J=2} I_t^j \left(\frac{\omega^{(j)}}{1 - \beta^{(j)}(L)} \right. \\
 &\quad \left. + \left[1 - \frac{(1 - \phi^{(j)}(L))(1 - (L))^{\delta^{(j)}}}{1 - \beta^{(j)}(L)} \right] \epsilon_t^2 \right), \\
 j &= 1 \text{ when } R_{oil,t-d} < r, \text{ and} \\
 j &= 2 \text{ when } R_{oil,t-d} \geq r.
 \end{aligned} \tag{17.1}$$

^{*} Summary statistics of the data for the 1983 to 1991 sub-period were also produced but are not reported. The Jarque–Bera and the Ljung–Box statistics in all sub-periods favour the application of a GARCH-type model.

[†] Sadorsky (1999) finds evidence that transmission of shocks from oil to stock returns is unidirectional. Hence, endogeneity of the latter can be assumed away.

[‡] Augmented Dickey–Fuller (Dickey and Fuller 1979, 1981) and the Phillips–Perron (Perron 1988, Phillips and Perron 1988) tests indicate that all returns are stationary. For the explanatory variables, the significance of the values pertaining to skewness (except for *HML*), kurtosis, J–B, Q(24), and all stationarity tests are similar to those of the sector returns. We also estimated the Threshold-GARCH (1, 1) specification for all sectors over the three sub-periods. Results show that, for an overwhelming number of cases, either shock persistence is very high, or returns follow an IGARCH process, highlighting the permanence of the shock effects on conditional variance. To study the possibility of permanence of shock effects on volatility, the more flexible FIGARCH model is appropriately suited.

TABLE 17.1 Summary Statistics

Series	Mean (%)	Std (%)	Min (%)	Max (%)	Skew	Kurtosis	J-B (p-value)	Q(24)	ADF(4)	ADF(4,t)	PP(0)	PP(4)	CV
Panel A: January 25 1991 to December 10, 1998													
Coal	0.147	1.830	-7.923	9.748	0.436	1.668	0.000	165.221	-22.651	-22.672	-58.482	-60.036	12.419
Oil-Gas Extraction	0.146	0.809	-5.734	4.750	-0.602	5.287	0.000	294.346	-15.814	-16.018	-34.905	-35.388	5.525
Building	0.173	1.374	-5.821	10.238	0.612	3.634	0.000	107.227	-18.236	-18.282	-44.409	-44.625	7.921
Chemical	0.127	0.837	-6.680	3.815	-1.217	7.705	0.000	332.129	-15.740	-15.814	-34.773	-35.293	6.602
Petroleum Refinery	0.072	0.860	-4.679	8.017	0.535	6.090	0.000	29.487	-18.909	-18.988	-44.253	-44.301	11.869
Plastic and Rubber	0.119	0.776	-4.020	6.352	0.325	4.963	0.000	83.406	-16.961	-17.181	-45.661	-45.764	6.520
Transport Equipment	0.119	0.736	-5.556	6.982	-0.295	8.496	0.000	100.207	-16.905	-17.011	-40.463	-40.833	6.202
Air Transportation	0.117	1.213	-6.558	6.295	0.154	2.548	0.000	50.600	-17.399	-17.425	-44.191	-44.269	10.345
Electric-Gas Services	0.081	0.459	-3.383	2.179	-0.433	3.293	0.000	47.679	-18.205	-18.201	-40.839	-41.000	5.693
Depository Inst.	0.176	0.513	-3.969	4.144	-0.654	9.386	0.000	752.827	-14.437	-15.045	-34.425	-35.242	2.907
ROIL	-0.035	1.921	-10.168	14.231	0.217	4.875	0.000	41.784	-20.766	-20.786	-43.649	-43.651	-54.759
RM	0.054	0.762	-6.620	4.810	-0.697	8.820	0.000	45.893	-20.535	-20.536	-41.509	-41.473	14.168
SMB	-0.008	0.469	-2.450	2.200	-0.136	1.417	0.000	81.534	-18.007	-18.227	-41.515	-41.609	-57.521
HML	0.016	0.423	-2.400	2.840	0.290	4.159	0.000	117.862	-19.288	-19.318	-36.308	-36.310	27.041
Panel B: December 11, 1998 to December 29, 2006													
Coal	0.116	1.522	-7.954	8.183	0.008	1.589	0.000	27.635	-19.532	-19.546	-42.971	-42.944	13.081
Oil-Gas Extraction	0.141	1.280	-5.242	5.187	-0.172	0.992	0.000	111.165	-18.598	-18.601	-36.707	-36.593	9.051
Building	0.134	1.628	-6.935	24.585	1.856	26.081	0.000	34.104	-19.573	-19.593	-42.386	-42.408	12.129
Chemical	0.118	1.196	-9.169	7.639	-0.208	4.260	0.000	247.060	-16.945	-16.995	-35.791	-36.044	10.120
Petroleum Refinery	0.121	1.175	-4.819	4.419	-0.099	0.678	0.000	32.314	-19.228	-19.237	-43.210	-43.198	9.713
Plastic and Rubber	0.087	0.931	-5.068	4.448	-0.057	1.804	0.000	85.329	-17.612	-17.609	-41.740	-42.007	10.648
Transport Equipment	0.069	0.950	-6.481	4.795	-0.241	2.719	0.000	85.432	-18.018	-18.013	-40.277	-40.576	13.750
Air Transportation	0.069	1.462	-19.275	6.985	-0.914	16.666	0.000	57.972	-19.269	-19.264	-40.955	-41.133	21.210
Electric-Gas Services	0.073	0.766	-6.278	3.544	-0.514	4.417	0.000	48.780	-19.545	-19.555	-41.916	-42.038	10.433
Depository Inst.	0.077	0.476	-2.461	2.550	-0.057	1.668	0.000	48.344	-18.313	-18.309	-42.904	-42.988	6.178
ROIL	0.086	2.346	-16.545	8.113	-0.588	2.954	0.000	23.235	-21.115	-21.136	-45.014	-45.030	27.408
RM	0.016	1.131	-6.650	5.310	0.092	2.367	0.000	44.497	-21.422	-21.431	-44.499	-44.505	71.546
SMB	0.025	0.630	-4.570	2.900	-0.578	3.697	0.000	88.119	-17.877	-17.884	-40.998	-41.152	25.300
HML	0.030	0.676	-4.930	3.900	-0.102	4.920	0.000	102.635	-20.294	-20.307	-40.005	-39.975	22.409

The critical values for ADF(4), ADF(4,t), PP(0) and PP(4) are 3.45, 3.96, 3.43 and 3.43, respectively at the 1% level. Q(24) is the Ljung-Box test for the 24th order serial correlation in return. The critical value for Q(24) at the 1% and 5% levels are 42.98 and 36.42, respectively. J-B is the Jarque-Bera joint normality test statistics.

In this model, $R_{i,t}$ is the excess return on sector i ($i = 1, 2, \dots, 10$), the market risk premium (RM), size (SMB) and book-to-market (HML) are the Fama–French factors, R_{oil} is the return on oil as defined previously, I_t^j is an indicator function for regime j ($j = 1$ and 2) and ε_t is the normal random error with mean of zero and variance σ^2 , conditional on the information set Φ_t available at time t . In the volatility equation, $\omega^{(j)}$ is a constant, $\phi^{(j)} \equiv [1 - \alpha^{(j)}(L) - \beta^{(j)}(L)](1 - L)^{-1}$ is the moving average (MA) component of the short-term volatility dynamics where $\alpha^{(j)}$ and $\beta^{(j)}$ are the ARCH and GARCH parameters, respectively, L denotes the lag operator, and, finally, δ is the fractional order of differencing, $0 < \delta < 1$, which accounts for long-term persistence in the FIGARCH specification. We describe the components of the model below.

17.2.2.1 The Return and Conditional Volatility Generating Processes

To estimate the industry returns, in excess of the risk-free rate, we propose a four-factor threshold model that adds return on oil (R_{oil}) to the market (RM), size (SMB) and book-to-market (HML) factors used by Fama and French (1992), to account for the oil price effects on sector return.* The mean equation specification of our model follows the generalization of Tsay (1998), which allows for exogenous threshold variables ($Z_{t-d} = R_{oil,t-d}$) and the simultaneous estimation of the delay parameter d and threshold variable r .

Li and Li (1996) and Brooks (2001) are among the few studies employing a Threshold GARCH-type model to address asymmetries in equity and currency returns, respectively. These studies use a Self-Exciting Threshold Autoregressive (SETAR)-GARCH specification to capture the asymmetries in equities return levels and variances.† Several studies have also found the prevalence of IGARCH properties in high-frequency financial data, implying indefinite persistence of the innovations in the conditional variance process. Poterba and Summers (1986) have also documented that the extent to which conditional volatility affects stock prices through a time-varying risk premium depends on the degree of shock persistence; transitory shocks to conditional volatility have smaller impacts on prices than long-term shocks. Thus, a precise depiction of the degree of volatility persistence is crucial to asset valuation. Indeed, neither the stationary GARCH nor the IGARCH specification has desirable features from an asset pricing point of view. Baillie *et al.* (1996) show that IGARCH properties observed in asset pricing data may be the trace of a mean-reverting long-memory FIGARCH data-generating process and propose the latter model as a way to allow for an intermediate degree of volatility persistence.*

The FIGARCH (1, δ , 1) model nests the GARCH (1, 1) and the IGARCH (1, 1) processes. When $\delta = 0$, FIGARCH becomes a GARCH process, with an exponential decay of the conditional variance. When $\delta = 1$, FIGARCH becomes IGARCH, where shocks to the conditional variance persist indefinitely. The FIGARCH (1, δ , 1) process is strictly stationary and ergodic if $0 \leq \delta \leq 1$ (Baillie *et al.* 1996). Indeed, since the conditional variance under the FIGARCH specification follows a hyperbolic rate of decay, it allows us to observe the true extent of long memory in the conditional variance of returns. Our model falls into

* To determine the appropriateness of the four-factor model used here versus a two-factor process (RM and R_{oil}), we also estimated a model excluding the other two factors. The log likelihood ratio test rejects the two-factor model for all sectors at 1%, clearly establishing the superiority of our four-factor specification.

† Unlike the threshold approach used here, where the regimes are determined by a structural indicator (R_{oil}), in a Self-Exciting Threshold Autoregressive model, the behaviour of a return series (R_t) changes once the series itself enters a different regime. The switch from one regime to the next is dependent on the past values of the series itself.

* Misspecification of volatility persistence has serious consequence on asset valuation. To assess this effect, we ran GARCH-M and FIGARCH-M models of the above specification. Using the coefficients of the volatility risk premium obtained from GARCH-M and FIGARCH-M models, a 10 basis points increase in conditional volatility leads to a \$7.13 billion and \$3.14 billion decline in the total market capitalization of the Oil and Gas industries, respectively. Detailed results can be obtained from the authors.

the class of threshold models transformed by Tsay (1989, 1998) and Li and Li (1996).^{*} It combines two piecewise-linear regimes to form a non-linear specification capable of modelling a wide array of empirical properties in asset returns, such as jumps, mean reversion and asymmetric and cyclical movements. The DT-FIGARCH model is a two-regime specification, each of which is piecewise linear with respect to oil returns $R_{oil,t-d}$. When the oil return is below (above) the threshold value (r), the trajectory of sector returns follows the first (second) regime. Thus, in regime 1 (regime 2) the coefficients of the model account for oil return effects that are below (above) the threshold value.

17.2.2.2 The Estimated Model

The mean and volatility equations described earlier are presented in more familiar terms below where all variables are as defined previously. The coefficient estimates for the model are obtained using the quasi-maximum likelihood method.[†]

$$R_{i,t} = \begin{cases} \left\{ \begin{aligned} &(b_0^1 + b_1^1 RM_t + b_2^1 SMB_t \\ &+ b_3^1 HML_t + b_4^1 R_{oil,t-d} + \varepsilon_t) \end{aligned} \right\} & \text{if } R_{oil,t-d} < r, \\ \left\{ \begin{aligned} &(b_0^2 + b_1^2 RM_t + b_2^2 SMB_t \\ &+ b_3^2 HML_t + b_4^2 R_{oil,t-d} + \varepsilon_t) \end{aligned} \right\} & \text{if } R_{oil,t-d} \geq r, \end{cases}$$

$$\sigma_t^2 = \begin{cases} \left\{ \begin{aligned} &\omega_0^1 + \beta_0^1 \sigma_{t-1}^2 + [1 - \beta_0^1(L)] \\ &[-(1 - \phi_0^1(L))(1 - (L))^{\delta^1}] \varepsilon_t^2 \end{aligned} \right\} & \text{if } R_{oil,t-d} < r, \\ \left\{ \begin{aligned} &\omega_0^2 + \beta_0^2 \sigma_{t-1}^2 + [1 - \beta_0^2(L)] \\ &[-(1 - \phi_0^2(L))(1 - (L))^{\delta^2}] \varepsilon_t^2 \end{aligned} \right\} & \text{if } R_{oil,t-d} \geq r. \end{cases}$$

Oil price effects may impact industry earnings via direct or indirect channels of influences. Specifically, oil price effects on industry returns may manifest themselves through the sensitivity of industry returns to oil prices as well as indirectly via the sensitivity of other risk factors to oil price movements. If oil price effects are subject to a threshold at the industry level, then it is likely that oil price changes above (below) the threshold trigger differential readjustments in how oil price and other risk factors are priced into industry returns. In the model described above, when the threshold variable (R_{oil}) exceeds the threshold value (r), it triggers a shift in the slopes of all the model factors with a delay window of (d) periods.^{*} We allow for shifts in all of the slopes for two reasons. First, this approach allows us to maintain generality and to test whether such shifts do occur, rather than assuming they do or they do not. Second, Fama and French (1997) argue that bad news about future cash flows tends to raise the industry's risk loading on *SMB* and *HML* factors. Accordingly, in our model, the increase in oil prices, if large enough, would tend to limit the cash flows of the oil-consuming sectors and to shift up the coefficients on the *SMB* and *HML* factors. Similarly, industries with large positive *HML* slopes (i.e. high BV/MV) are likely to have

^{*} The advantages of the Tsay (1998) threshold method are threefold. First, this method is not an exclusively self-exciting threshold autoregressive type model, where the threshold value is determined by its own returns, it rather allows for exogenous variables to be considered for threshold. Second, it allows for the estimation of the delay parameter (d) and the threshold value (r) simultaneously, using the Akaike Information Criteria, instead of scatter plots of predictive residuals used by Tsay (1989) that require subjective assessment. Third, the test procedure to detect change points is relatively simple and has an asymptotic Chi-squared distribution.

[†] Appendix A presents a discussion of the simultaneous estimation of the threshold coefficient (r) and the delay parameter (d) for each sector.

^{*} The issue of how many thresholds should be estimated and the economic interpretation of the threshold values are still unresolved in the extant literature. Meyer and von Carmon-Taubadel (2004) provide an extensive review of the asymmetric price transmission literature and conclude that the magnitude of threshold values have never been interpreted in an economic sense. The estimated values are generally interpreted as the minimum incentives required by the economic agent to elicit price or return adjustments.

experienced surprise negative conditions (or to be relatively more distressed) and tend to sell low (Fama and French 1997). Along the same lines, the *HML* slopes would tend to shift higher (lower) with higher oil prices for the oil-consuming (substitute) sectors. In brief, the Fama–French argument implies that when oil return increases are above a positive threshold ($j = 2$), the coefficients (b_2^2, b_3^2), which measure the sectoral effects of *SMB* and *HML* for the oil-consuming sectors, will be larger in magnitude than their counterparts (b_2^1, b_3^1) when oil return increases are below the threshold ($j = 1$) because they have to spend more on their oil consumption than in the former regime. To test this view, we allow for a shift in the slope of these factors.*

17.2.3 Testable Hypotheses

Two main sets of hypotheses are tested. First, the appropriateness of the FIGARCH specification, and second the asymmetry of the effects of the model factors (R_{oil} , *RM*, *SMB*, *HML*) under the two regimes ($J = 1, 2$). Correct specification of the stock return patterns has major implications on the empirical findings and related policy conclusions. For example, since the systematic risk measure (beta) of a sector is used to estimate the cost of capital for that sector, its accurate measurement is essential for deriving a correct estimate of the cost of capital and making optimal investment decisions. In particular, if projects on alternative energy sources are assessed based on incorrect cost of capital estimates, the choices made will be unreliable. Similarly, accurate measurement of sectoral sensitivities to oil price changes can serve as a basis for policies concerning energy source diversity (e.g. development of renewable energy such as fossil fuel, atomic energy, etc.) and can help determine optimal hedging positions.

Finding of dissimilar (asymmetric) effects under the two regimes is also important because it may necessitate differential regulatory policies by the government, and dissimilar investment and managerial decisions under these regimes. In particular, if the effects of oil price changes on the stock returns are symmetric, the impacts of oil price increases and oil price decreases (bad news and good news) would be similar in magnitude and counterbalancing. However, if oil price increases harm a sector more than oil price declines improve its conditions, firms in the sector would need to take steps to hedge the shocks, or to be prepared to absorb them. Investors would also have to take this pattern of behaviour into account when investing in the sector and the government may consider subsidies to moderate the adverse effects from the shocks.

The hypotheses concerning the functional form are as follows.

H_{11} : The proper specification for sector returns is GARCH (1, 1) (the effect of a shock decays at an exponential rate): $H_0 : \delta^1 = \delta^2 = 0$

H_{12} : The proper specification for sector returns is IGARCH (shocks persist indefinitely): $H_0 : \delta^1 = \delta^2 = 1$

H_{13} : The proper specification for sector returns is the Fama–French model: $H_0 : b_4^1 = b_4^2 = \beta_0^1 = \beta_0^2 = \phi_0^1 = \phi_0^2 = \delta^1 = \delta^2 = 0$

Asymmetry hypotheses are given below. In each of the first four cases, symmetry indicates that sensitivity of the sectoral return to the market, *SMB*, *HML*, or R_{oil} is the same, regardless of whether the changes in these variables occur when oil returns are above or below the threshold (r). The last hypothesis, H_{25} , investigates whether the volatility decay process is symmetric.

H_{21} : Market risk premia (slopes) are symmetric for oil return shocks above and below the threshold: $H_0 : b_1^1 = b_1^2$

H_{22} : *SMB* risk premia (slopes) are symmetric for oil return shocks above and below the threshold: $H_0 : b_2^1 = b_2^2$

H_{23} : *HML* risk premia (slopes) are symmetric for shocks above and below the threshold $H_0 : b_3^1 = b_3^2$

* An alternative would be to introduce a shift only in the slope of the oil return variable (R_{oil}). This would likely produce a bigger shift in the slope of R_{oil} . Currently, the effects are divided among all four factors. Hence, our specification may underestimate the shift in the slope of R_{oil} .

H_{24} : Oil return risk premia (slopes) are symmetric for oil return shocks above and below the threshold: $H_0 : b_4^1 = b_4^2$

H_{25} : Volatility decay process is symmetric for oil price shocks above and below the threshold: $H_0 : \delta^1 = \delta^2$

17.3 Empirical Findings*

17.3.1 Tests of Model Specification

It is crucial to determine the precise measure of volatility persistence (memory) because of its implications for asset valuation. A FIGARCH specification is employed here because it nests the GARCH and IGARCH models and it also allows us to observe the true extent of long memory in the conditional variance of returns. Another major issue is whether oil shock effects show similar decay patterns when oil returns are below and above the threshold because differential decay patterns and different decay speeds call for different policy actions by the oil producers, oil consumers, and government regulators.

Results for model specification tests are reported in Table 17.2. Based on the statistics reported in this table, the basic GARCH, IGARCH and the Fama–French specifications (H_{11} , H_{12} , H_{13}) are rejected for almost all sectors in both the volatile and calm oil price periods (Petroleum Refinery is the exception in the calm period). The implication is that these specifications are indeed too restrictive to capture the behavioural patterns of sectoral stock returns and their application may produce distorted results.[†] The rejection of both GARCH and IGARCH specifications also indicates that oil price shock effects do not decline exponentially (geometrically), nor persist indefinitely. Instead, these effects decay at a slower hyperbolic pace, reflecting a long-memory pattern. In addition, the rejection of the Fama–French model indicates that the menu of factors chosen by these authors is too restrictive, and other factors, at least the oil returns (R_{oil}), are likely to play a role in describing asset returns. Overall, these findings indicate that predictions and policy recommendations based on the functional forms rejected by the data are misleading and can fall short of or run counter to their intended objectives.

17.3.2 Tests of Asymmetry

An interesting question is whether the systematic risks associated with the Fama–French and oil return factors included in the model change when oil return shocks become large enough to exceed the corresponding thresholds. The answer to this question would help us reach a more accurate assessment of the overall sectoral risk and has implications for investors and firm managers. Hypotheses H_{21} , H_{22} , H_{23} and H_{24} investigate these symmetry issues, namely whether the slopes of the model factors are identical when oil returns fall short and exceed the threshold values. A second question is whether oil shock effects show similar decay patterns when oil returns are below and above the threshold, because differential decay patterns and different decay speeds call for different policy actions by the oil producers, oil

* Diagnostic statistics for each model including the Log Likelihood value, the LjungBox portmanteau test for up to 20th order serial correlation in the standardized residuals $Q(20)$ and squared standardized residuals $Q^2(20)$ are presented at the bottom of Tables 17.3(A) and 3(B). The $Q^2(20)$ statistics suggest that the FIGARCH model adequately captures the autocorrelations in the conditional variance of sector returns.

[†] The magnitude of the distortion due to model misspecification can be considerable. We contrasted the oil return effects on the Oil–Gas Extraction sector derived from the single-regime GARCH and the model used here for 1998–2006. The differences are statistically significant and sizeable. Under GARCH, the ROIL coefficient was 0.21 ($t = 31.8$) while under FIGARCH the ROIL coefficient was 0.19 ($t = 70.1$) above the threshold, and 0.21 ($t = 79.35$) below it, indicating an overestimation of two basis points under GARCH. With a market capitalization of Oil–Gas Extraction of \$793 008 million in 2006 for firms with SIC 13, a two basis point overestimation translates into approximately (\$793 008 million $\times$ two basis points) or \$158.60 million.

TABLE 17.2 Hypotheses Concerning Functional Form and Symmetry

Hypotheses Concerning Functional Form						Symmetry Hypotheses					
H ₁₁ : The proper specification is GARCH(1, 1): $\delta^1 = \delta^2 = 0$						H ₂₁ : Market risk premiums are symmetric: $b_1^1 = b_1^2$					
H ₁₂ : The proper specification is GARCH: $\delta^1 = \delta^2 = 1$						H ₂₂ : SMB premiums are symmetric: $b_2^1 = b_2^2$					
H ₁₃ : The proper specification is Fama–French: $b_4^1 = b_4^2 = \beta_0^1 = \beta_0^2 = \Phi_0^1 = \Phi_0^2 = \delta^1 = \delta^2 = 0$						H ₂₃ : HML premiums are symmetric: $b_3^1 = b_3^2$					
						H ₂₄ : Oil price effects are symmetric: $b_4^1 = b_4^2$					
						H ₂₅ : Volatility decay process is symmetric: $\delta^1 = \delta^2$					
			Electric-Gas Services (SIC49)	Oil-Gas Extraction (SIC13)	Petroleum Refinery (SIC29)	Building (SIC15)	Air Transportation (SIC45)	Chemical (SIC28)	Plastic and Rubber (SIC30)	Transport Equipment (SIC37)	Depository Inst. (SIC60)
	Coal (SIC12)										
January 25, 1991 to December 10, 1998 Sub-Period											
H ₀ : Model Specification											
H ₁₁ :	2d.f.	879.47	1377.08	306.22	3.23E-12	901.31	2286.05	248.94	154.63	431.51	3991.12
H ₁₂ :	2d.f.	1919.5	2674.51	1521.22	3486.72	4582.92	57648.78	2604.19	2107.99	2757.24	1618.66
H ₁₃ :	8d.f.	7.76E+04	2.37E+05	3.40E+04	4.26E+04	2.73E+05	1.75E+07	6.62E+04	4.26E+04	2.73E+05	1.97E+05
H ₀ : Symmetry											
H ₂₁ :	1d.f.	2.15	12.04	0.03	18.15	0.06	3.46	535.44	5.55	23.76	89.44
H ₂₂ :	1d.f.	18.3	126.99	179.08	2.18	45.28	72.58	501.15	6.3	85.34	10.16
H ₂₃ :	1d.f.	1.26E-03	80.4	30.52	1.45	2.17	23.38	104.6	27.07	44.63	0.02
H ₂₄ :	1d.f.	1.94	3.6	0.4	0.3	0.65	0.73	5.59E-05	1.23	2.79	1.91
H ₂₅ :	1d.f.	189.41	298.88	2.41	2.88E-12	24.62	78.76	108.2	32.36	289.87	411.92
December 11, 1998 to December 29, 2006 Sub-Period											
H ₀ : Model Specification											
H ₁₁ :	2d.f.	1153.09	1500.75	506.87	119.68	17965.15	302.58	1143.79	79.69	1476.86	351.96
H ₁₂ :	2d.f.	1458.12	8871.43	3118.18	1871.46	8.76E+06	2622.21	1497.71	2473.21	1881.27	1562.25
H ₁₃ :	8d.f.	6.76E+05	1.44E+05	1.92E+05	7.05E+04	2.05E+04	2.94E+03	5.04E+05	3.23E+05	2.51E+05	6.97E+04
H ₀ : Symmetry											
H ₂₁ :	1d.f.	104.47	593.95	20.97	9.81	79.86	1586.94	621.56	192	119.28	196.57
H ₂₂ :	1d.f.	59.29	28.32	0.13	41.95	17.02	5.29	16.08	45.71	98.87	13.8
H ₂₃ :	1d.f.	6.73	180.98	8.99	40.86	472.5	205.58	40.59	4.95	646.46	751.66
H ₂₄ :	1d.f.	0.8	3.15	5.55	1.2	4.74	0.04	6.73	17.4	8.49	0.61
H ₂₅ :	1d.f.	66.76	105.15	146.71	3.54	104.73	79	67.33	0.19	456.62	8.85

The critical values of $\chi^2(2d.f) = 5.9, 4.6, \chi^2(1d.f) = 3.8, 2.7$ and $\chi^2(8d.f) = 15.5, 13.3$ at the 5% and 10% levels, respectively.

consumers and government regulators. Hypothesis H_{25} ($H_0: \delta^1 = \delta^2$) tests whether the speed of memory decay is identical for smaller oil returns falling below the threshold and the larger ones positioned above it. The results of symmetry tests can provide an answer to these questions.

According to the test results in Table 17.2, in the volatile period (1998–2006), the symmetry of the response to market (*RM*) and *HML* factors is unanimously rejected for all 10 sectors, and for all, except for the Oil–Gas Extraction sector, for the *SMB* factor. Symmetry results are less uniform in the low volatility period (1991–1998). Nevertheless, symmetry is always rejected in more than two-thirds of the cases for each of the three Fama–French factors. These findings indicate that the systemic risks (the slopes) of the three Fama–French factors do indeed change when the oil return passes the threshold level during both the calm and volatile periods, although the shift (regime change) occurs in more sectors when oil prices are more volatile and larger overall. We also note that the direction of the shift is different for different sectors. In particular, in both periods, the market risk for the oil-substitute sector, Coal, falls while that for the oil-consuming sector, Air Transportation, increases, when oil return exceeds the threshold, indicating that the former benefits while the latter suffers (becomes riskier) when oil returns exceed the threshold. This pattern is observed over a broader spectrum of industries in the less volatile period.

These findings provide only mixed support for the Fama–French (1997) argument that bad news about future cash flows tends to raise the industry's risk loading on *SMB* and *HML*. The bad news here for the oil-consuming sector is the increase in the oil return above the threshold level. For the Air Transportation and Transport Equipment sectors, the increase in oil return does, as expected, shift up the *SMB* and *HML* risk in the calm period, but it fails to do so in the volatile period. Similarly, for the oil-substitute sector, Coal, the results are mixed; an increase in oil return does reduce the *SMB* and *HML* risk in the volatile period, but not in the calm period. This indicates that only when markets are volatile do oil price changes translate into better conditions for the oil-substitute sector. These findings provide considerable evidence against the extant studies that assume identical coefficients for the risk factors in their models at all levels of oil return change as a part of their maintained hypotheses.

It is notable that the shift in the slope (systematic risk) of the fourth factor, the oil return (R_{oil}), is less frequent in both sub-periods, and more so in the calm periods, compared with the other three factors. This is perhaps an indication that oil return, although a significant factor, plays a less important role in determining industry returns than the traditional Fama–French factors. This issue will be further discussed in Section 17.3.3.3 when analysing the economic significance of the model factors. A stronger explanation may be that the effects of increases in oil return are mostly picked up by changes in the oil return level (R_{oil}) itself, rather than by a change in its coefficient. In addition, it is likely that the Fama–French factors, e.g. the market return, also proxy for oil price changes. It follows that if the interdependence of the oil and market returns is stronger when oil returns are higher (above threshold), market return may pick up some of the asymmetry in response to the oil return and weaken the asymmetry test for the latter variable.

Based on the slope shifts in the four factors, two conclusions can be drawn. First, the effect of oil price changes includes both direct and indirect components. Thus, the overall effect should be assessed by adding up the effects due to the slope shifts of all four factors in the model, including oil return and Fama–French factors. If the indirect effects are not accounted for in government policy decisions and managerial planning by firms, the resulting cost/benefits analysis of oil price changes will be distorted and, hence, the adopted policy decisions will be suboptimal. Second, findings based on simpler models are suspect as they fail to account for the asymmetry of sectoral responses to changes in the model factors when oil returns are below and above the threshold.

The hypothesis of identical speeds of decay in volatility when oil returns are below and above the threshold (H_{25}) is rejected strongly for all sectors, except for Oil–Gas Extraction and Petroleum Refinery sectors in the calm period, and the Plastic and Rubber sector in the volatile period. The results of this hypothesis indicate that the volatility decay process is indeed dissimilar for oil return changes below and above the threshold and, hence, these dissimilarities have to be accounted for. The estimated values of the decay parameter—, reported in Tables 17.3(A) and 3(B), are found to be, with very few exceptions,

TABLE 17.3A Double-Threshold FIGARCH(1, δ , 1) Results January 25, 1991 to December 10, 1998 Sub-Period

1	2	3	4	5	6	7	8	9	10	11	12
Coefficient	Variable	Coal (SIC12)	Electric- Gas Services (SIC49)	Oil-Gas Extraction (SIC13)	Petroleum Refinery (SIC29)	Building (SIC15)	Air Transportation (SIC45)	Chemical (SIC28)	Plastic and Rubber (SIC30)	Transport Equipment (SIC37)	Depository Inst. (SIC60)
b_0^2	CONSTANT	1.35E-03	2.34E-04	1.45E-03	-1.40E-03	1.90E-03	1.46E-04	5.72E-04	1.93E-04	1.39E-04	1.63E-03
		2.82	10.593	14.833	-22.80	7.872	0.797	23.628	3.487	5.744	33.396
b_0^1	CONSTANT	-1.21E-04	5.26E-05	4.41E-04	-8.95E-04	1.79E-04	1.85E-04	5.16E-04	8.29E-04	7.00E-04	1.05E-03
		-1.074	1.442	15.252	-15.259	2.739	4.993	18.924	16.288	11.668	53.695
b_1^2	RM	0.785	0.597	0.839	0.685	1.1	1.192	1.002	0.714	0.95	0.712
		10.975	204.311	96.893	104.036	20.01	40.255	489.536	84.894	218.994	126.141
b_1^1	RM	0.893	0.576	0.837	0.733	1.086	1.135	0.892	0.688	0.895	0.658
		47.076	110.593	197.539	57.111	260.329	145.161	208.285	90.632	85.772	496.254
b_2^2	SMB	0.126	0.31	0.509	0.332	1.344	1.199	0.921	0.685	0.728	0.544
		1.213	93.814	30.217	20.989	24.021	23.627	224.24	57.693	191.481	31.813
b_2^1	SMB	0.586	0.222	0.768	0.266	0.946	0.75	0.725	0.644	0.601	0.488
		19.76	31.25	80.58	20.888	48.38	53.141	93.666	59.628	45.459	161.814
b_3^2	HML	0.411	0.328	0.256	0.536	0.557	0.634	-0.015	0.22	0.475	0.485
		3.189	63.872	10.412	29.796	5.144	11.339	-5.024	18.581	78.817	32.605
b_3^1	HML	0.406	0.418	0.395	0.471	0.717	0.356	0.035	0.308	0.364	0.483
		11.804	48.895	72.9	23.922	63.413	25.782	9.026	25.277	23.497	190.624
b_4^2	R_{oil}	-0.034	0.003	0.071	0.04	0.043	-0.05	-0.007	0.005	-0.005	0.009
		-0.995	1.961	5.735	5.908	1.222	-1.747	-6.807	0.797	-2.968	0.876
b_4^1	R_{oil}	0.015	-0.01	0.063	0.035	0.015	-0.025	-0.007	-0.004	0.008	-0.005
		1.803	-1.463	56.361	4.795	3.802	-6.489	-0.488	-0.774	1.07	-3.424
ω_0^2	CON. VOL.	-1.86E-13	-7.27E-12	-2.43E-14	1.84E-13	0.00E+00	0.00E+00	0.00E+00	-1.13E-14	-1.77E-14	-8.91E-08
		-2.53E-08	-7.40E-05	-4.12E-07	0.023	9.40E-10	8.96E-06	-2.42E-07	-4.67E-07	-0.015	-0.226
ω_0^1	CON. VOL.	1.14E-06	9.96E-14	6.20E-07	8.73E-15	0.00E+00	0.00E+00	1.55E-07	-6.22E-13	4.52E-13	4.87E-16
		0.734	3.11E-07	64.374	9.50E-07	-1.72E-07	0.00E+00	0.262	-0.001	5.38E-05	1.81E-08
β_0^2	GARCH	0.581	0.636	0.231	0.944	0.949	1.332	0.144	0.375	0.418	0.364
		15.895	37.784	6.521	267.238	23.642	129.905	5.937	10.624	15.993	17.536

β_0^1	GARCH	0.849	0.848	0.507	0.561	0.777	0.753	0.575	0.271	0.77	0.97
		64.787	28.632	22.981	27.865	68.77	117.436	12.712	9.486	23.999	112.454
Φ_0^2	PHI	0.748	0.484	-3.86E-08	0.887	1.045	1.285	0.107	0.3979	0.4453	0.216
		23.288	25.983	0	180.745	36.131	135.535	4.476	11.676	17.598	9.34
Φ_0^1	PHI	0.458	0.131	0.232	0.736	0.613	0.624	-5.00E-09	5.13E-07	2.18E-08	0.112
		20.065	5.606	9.255	43.126	40.318	77.311	-1.01E-07	1.82E-05	5.02E-07	6.046
δ^2	DECAY	0.019	0.198	0.267	0.129	0.123	0.078	0.159	0.1015	0.0839	0.351
		0.608	11.635	8.304	9.845	3.229	7.629	9.654	4.203	4.801	20.073
δ^1	DECAY	0.582	0.832	0.323	0	0.341	0.223	0.67	0.279	0.999	0.951
		28.482	30.032	16.302	0	26.991	29.682	13.751	12.136	19.967	50.299
LL Value		5335.09	8770.41	7502.96	7068.489	6170.7	6457.57	8198.49	7377.32	7982.39	8785.82
Q(20)		174.09	48.56	234.58	52.03	63.89	28.62	212.26	89.76	64.39	368.33
Q ² (20)		16.09	11.16	7.05	33.72	13.72	10.5	9.22	8.31	16.64	11.82
Threshold (%)		1.51	-0.687	0.977	-0.477	1.621	1.158	-1.447	-0.309	-0.805	1.588
Delay (days)		4	3	1	7	6	7	3	2	4	7

t-values are below the coefficients. Significant *t*-values at the 5% level are highlighted in bold. The critical *t*-values are 1.64, 1.96, 2.57 and Q(20) are 28.41, 31.41 and 37.57 at the 10%, 5% and 1% levels, respectively.

TABLE 17.3B Double-Threshold FIGARCH(1, δ , 1) Results December 11, 1998 to December 29, 2006 Sub-Period

1	2	3	4	5	6	7	8	9	10	11	12
Coefficient	Variable	Coal (SIC12)	Electric-Gas Services (SIC49)	Oil-Gas Extraction (SIC13)	Petroleum Refinery (SIC29)	Building (SIC15)	Air Transportation (SIC45)	Chemical (SIC28)	Plastic and Rubber (SIC30)	Transport Equipment (SIC37)	Depository Inst. (SIC60)
b_0^2	CONSTANT	9.01E-04	3.31E-04	6.27E-04	5.62E-04	8.59E-04	6.98E-05	5.50E-04	2.37E-04	2.16E-04	2.66E-04
		6.406	8.997	7.742	4.263	19.700	0.434	12.951	4.040	3.646	19.163
b_0^1	CONSTANT	-1.65E-04	4.93E-05	2.34E-04	2.88E-04	3.05E-04	-2.47E-04	8.19E-04	2.57E-04	-1.03E-04	6.80E-04
		-3.78	1.515	3.673	4.081	3.953	-4.997	11.652	5.049	-2.258	25.16
b_1^2	RM	0.738	0.649	0.89	0.899	0.955	1.342	0.782	0.722	0.798	0.445
		76.256	305.544	103.527	58.976	425.55	217.662	245.428	126.086	133.642	377.24
b_1^1	RM	0.858	0.743	0.789	0.849	1.045	1.02	0.931	0.624	0.872	0.409
		130.617	232.156	215.166	166.525	106.947	196.423	183.095	151.569	266.207	178.731
b_2^2	SMB	0.449	0.193	0.53	0.204	0.615	0.61	0.762	0.369	0.366	0.241
		19.654	31.219	28.803	10.254	109.342	22.947	95.699	38.383	40.171	86.004
b_2^1	SMB	0.636	0.241	0.526	0.35	0.667	0.543	0.815	0.448	0.475	0.253
		76.882	36.434	45.741	33.788	59.247	42.92	77.527	68.659	78.083	152.787
b_3^2	HML	0.773	0.693	0.973	0.992	0.351	0.955	0.132	0.425	0.436	0.361
		17.437	204.501	89.478	55.392	78.13	45.652	29.197	46.365	47.101	177.98
b_3^1	HML	0.889	0.777	0.858	0.864	0.808	0.636	0.209	0.448	0.708	0.274
		156.494	148.447	92.209	97.616	39.325	82.526	18.53	82.031	130.007	111.954
b_4^2	R_{oil}	0.103	0.002	0.198	0.106	-0.075	-0.088	-0.014	-0.029	-0.001	-0.015
		4.626	0.453	70.192	4.346	-2.721	-4.452	-5.068	-3.561	-0.168	-8.742
b_4^1	R_{oil}	0.083	0.011	0.216	0.133	-0.014	-0.092	0.003	6.67E-03	-0.017	-0.013
		25.888	4.253	79.35	47.225	-2.807	-35.693	0.517	2.541	-9.661	-6.926
ω_0^2	CON. VOL.	-3.10E-13	0.00E+00	-8.14E-12	-9.58E-13	-8.32E-11	7.55E-06	5.67E-14	8.70E-07	-4.93E-15	-4.12E-14
		-1.48E-07	-7.43E-06	-4.53E-06	-2.64E-07	-1.67E-04	1.43	2.66E-05	0.957	-1.07E-04	-7.00E-07
ω_0^1	CON. VOL.	-2.67E-07	0.00E+00	1.21E-06	1.82E-13	2.79E-11	-4.39E-15	-9.45E-14	-2.06E-14	-1.61E-07	-2.59E-14
		-0.149	2.04E-08	0.895	1.46E-07	0	0	-0.029	0	-24.578	-1.96E-07

β_0^2	GARCH	0.658	0.991	0.511	0.166	0.887	2.61E-01	0.769	0.525	0.387	0.707
		26.553	42.91	18.339	4.541	396.468	5.555	46.169	17.32	15.126	33.214
β_0^1	GARCH	0.811	0.612	0.7	0.498	0.943	0.51	0.773	0.408	0.897	0.394
		56.277	43.45	35.248	21.334	135.009	22.53	32.974	16.715	42.431	12.327
Φ_0^2	PHI	0.446	1.039	0.464	0	0.833	0.232	0.526	0.448	0.28	0.513
		16.891	66.233	16.231	0	334.108	5.095	25.081	14.488	10.519	21.791
Φ_0^1	PHI	0.411	0.299	0.463	0.369	0.995	0.275	0.289	0.387	0.103	0.15
		19.758	19.274	20.473	15.262	212.56	10.675	10.92	15.81	4.517	4.648
δ^2	DECAY	0.333	0.057	0.148	0.149	0.069	0.024	0.382	0.161	0.177	0.283
		15.214	2.437	5.959	4.513	27.078	0.87	21.933	6.103	9.095	13.391
δ^1	DECAY	0.57	0.385	0.283	0.218	0.013	0.304	0.641	0.148	0.919	0.383
		30.283	31.115	13.597	10.388	4.342	17.388	24.871	7.506	35.492	13.953
LL Value		6095.84	8281.94	6568.97	6840.2	6119.08	6518.31	7582.38	7290.29	7825.51	9344.29
Q(20)		15.67	31.57	52.49	33.7	63.89	21.708	198.93	37.196	75.69	275.48
Q ² (20)		12.68	26.63	11.07	19.4	13.72	17.507	17.05	14.685	8.77	8.28
Threshold (%)		1.13	0.951	0.621	2.104	1.756	2.311	-0.541	0.795	0.499	-0.562
Delay (days)		7	1	1	6	5	7	6	7	2	4

t-values are below the coefficients. Significant *t*-values at the 5% level are highlighted in bold. The critical *t*-values are 1.64, 1.96, 2.57 and Q(20) are 28.41, 31.41 and 37.57 at the 10%, 5% and 1% levels, respectively.

smaller for oil returns positioned above the threshold in both periods, indicating a faster decay (less persistence) when return increases are larger. This may happen because unusually large oil price increases are more likely to be considered transitory and, hence, their effects die down more quickly. The dissimilarity of the volatility decay process is also overlooked in the literature.

17.3.3 Sector Analysis*

The estimation results for the 1991–1998 (calm) and 1998–2006 (volatile) periods are reported in [Tables 17.3\(A\)](#) and [3\(B\)](#), respectively.[†] In the 1991–1998 period, the threshold values (r) are positive for sectors such as Coal, Oil–Gas Services, Building, Air Transport and Depository Institutions, indicating that, for these sectors, regime shifts or changes in sensitivities occur when oil returns are rising. Among these industries, the Oil–Gas Extraction industry exhibits the lowest threshold level (0.977%) and Building and Depository Institutions have the largest threshold values (1.621%, 1.588%). These figures indicate that although it does take a considerable change in the return on oil to trigger a regime shift for either of these industries, the requirement is much smaller for the former than for the latter industries. It is notable that the remaining industries have negative thresholds, indicating that regime shifts or changes in factor sensitivities occur when oil return is falling. The magnitudes again vary across sectors.

17.3.3.1 Analysis of Oil Effects (R_{oil} Coefficient) Below and Above the Threshold

Assessing the effect of R_{oil} on sector return is of particular interest because of the major macroeconomic implications of oil price changes, especially given the spectacular rise in oil prices and the tremendous oil price volatility during the recent decades. The figures in [Tables 17.3\(A\)](#) and [3\(B\)](#) delineate the relationship between sectoral returns and R_{oil} . The direct marginal effect of oil price increases on sector return is determined by the magnitude of the R_{oil} coefficient b_4^2 when R_{oil} equals or exceeds the threshold (r), and by b_4^1 when it falls short of it. We analyse the results according to the four groups of industries considered and do so selectively within each group to save space.

17.3.3.1.1 Oil-Substitute Sectors

The results for the Coal sector, an oil-substitute, are presented in column 3 of [Tables 17.3\(A\)](#) and [3\(B\)](#). The threshold and the coefficients for this industry are $r = 1.51\%$, $b_4^1 = 0.015$ and $b_4^2 = -0.034$ (insignificant) in the calm period, and $r = 1.13\%$, $b_4^1 = 0.083$ and $b_4^2 = 0.103$ in the volatile period. These figures indicate that, in the volatile period, the increase in oil return needed for a regime change is smaller ($r = 1.13\%$ instead of 1.51%) and the oil effect on the Coal sector is much stronger, compared with the calm period, both when R_{oil} is below the threshold ($b_4^1 = 0.083$ versus $b_4^1 = 0.015$) and above it ($b_4^2 = 0.103$ versus b_4^2 insignificant). Specifically, in the more volatile period, the increase in sector return due to a one percentage point change in the oil return is 8.3 basis points (bp) compared with 1.5 bp during the calm period, when the increase in R_{oil} falls short of the threshold, and 10.3 bp compared with zero when R_{oil} exceeds the threshold.

The Coal sector does show a strong substitution effect during the period of oil price volatility as higher oil returns lead to greater coal sector returns. The coefficient of R_{oil} in this period is larger and more stringently significant compared with the calm oil price period, reflecting the greater sensitivity of

* As explained in the appendix, threshold values can be positive or negative and can be larger or smaller during periods of high oil price volatility compared to calm periods.

[†] The model was also estimated for the 1983–1991 period. Results are rather similar to the 1998–2006 period. A number of alternative specifications of Threshold-GARCH models were estimated as well. These include GARCH (1, 1), GARCH (1, 1)-M, and GARCH (1, 1) with market volume in the variance equation. Regardless of the GARCH specification used, the persistence of shocks (the sum of the ARCH and GARCH parameters), is 1.0 or close to 1.0 for most sectors, rendering the inferences unreliable. We also estimated FIGARCH (1, δ , 0) models. For space limitation, only the results of FIGARCH (1, δ , 1) are presented here. Other results can be obtained from the authors.

market participants under volatile conditions when oil price changes are more difficult to predict. This effect is slightly stronger when oil return surpasses the threshold. During the calm period, substitution is weaker, and when oil return surpasses the threshold, it loses significance. The lower strength of the relationship between oil return and risk premia when the return exceeds the threshold, in both periods, may indicate that when the oil return rises beyond a certain level, oil users are motivated to economize further on the use of energy and to switch to alternative energy sources such as solar power and wind. Moreover, protective hedges against oil price increases are more likely to be triggered under higher oil prices.

Electric–Gas Services may also be considered an oil substitute sector, although it is an oil consumer as well. For this industry, the threshold and the coefficients are $r = -0.68\%$, $b_4^1 = -0.010$ (insignificant), $b_4^2 = 0.003$, and $r = 0.95\%$, $b_4^1 = 0.011$, $b_4^2 = 0.002$ (insignificant) (column 4). The effect of oil return on the Electric–Gas Services sector is positive when statistically significant, as expected from an oil-substitute sector. Significant cases occur when oil returns rise above the threshold in the clam period and below it in the volatile period. Excessively high oil prices in the volatile period may have discouraged demand for this industry along with that for oil, in favour of other energy sources, and may have also made the passing on of the additional costs to the consumers difficult, offsetting any gains due to the substitution effect.

17.3.3.1.2 Oil-Related Sectors

The results for the oil-related industries, namely the Oil–Gas Extraction and Petroleum Refinery sectors, are reported in columns 5 and 6, respectively. The threshold and the oil return coefficients for the Oil–Gas Extraction sector in the calm and volatile periods are, respectively, $r = 0.98\%$, $b_4^1 = 0.06$, $b_4^2 = 0.07$ and $r = 0.62\%$, $b_4^1 = 0.216$, $b_4^2 = 0.198$. According to these figures, the effect of oil return increases on this sector is always positive, as expected from an oil-related industry. As the oil return increases, this industry experiences higher demand and higher prices, resulting in increased profitability, that in turn results in higher stock returns. In addition, as observed for the Coal industry, the threshold for the volatile period is smaller, indicating that it takes less of an increase in oil return to trigger a regime change. This may be because changes in the calm atmosphere of the earlier sub-period are more likely to be considered temporary, requiring a bigger oil price increase for a regime shift, or perhaps because, under volatile conditions, market participants are more accepting of a regime change. Another possibility is that the prevailing expectation formation mechanism was one of mean reversion such that market players expected a reversal in price trends, unless a considerably large increase occurred, convincing the market participants that it was time for a regime shift. As for a comparison between the two periods, the marginal effect of oil return on the Oil–Gas Extraction sector return is much higher in the volatile period than in the calm period (21.6 bp to 19.8 bp, compared with 6 bp to 7 bp in the clam period). For the Petroleum Refinery sector (column 6), the threshold and the oil return coefficients for the calm and volatile periods are, respectively, $r = -0.48\%$, $b_4^1 = 0.035$, $b_4^2 = 0.040$ and $r = 2.10\%$, $b_4^1 = 0.133$, $b_4^2 = 0.106$. This indicates that when oil return rises, the sector will benefit in all cases, but much more so during the volatile period, although these effects require a larger oil price increase to exceed the threshold. These findings are by and large consistent with expectations for an oil-related industry.

17.3.3.1.3 Oil-Consuming Sectors

This group of industries includes Building Construction, Air Transportation, Chemical, Plastic and Rubber and Transport Equipment (columns 7–11). The expectation may be that oil price increases will harm the stock returns of the firms in the oil-consuming industries. However, market conditions and market power may alter these effects. For example, if an industry has sufficient market power to pass on the additional costs due to oil price increases, or even more than the additional costs incurred, to the consumers, it may benefit from higher oil prices or, at least, it may remain unscathed. Similarly, if an industry is hedged against oil price increases, it need not be affected negatively. In a hedged scenario, the

pattern of the effect will depend on the conditions under which the hedge becomes effective, how close it is to a complete hedge, and what the partial hedge provides.*

Moreover, while the direct effect from oil price increases on a given industry may be negative, the industry may be subject to indirect positive effects, counterbalancing the direct effects. For example, if oil price increases lead to portfolio reshuffling by investors in favour of a particular oil-using industry, this industry's stock return may advance with higher oil prices, at least partially offsetting the direct impact.[†] Similarly, if sales contracts are indexed to oil prices, or if higher oil prices lead to supply side inflation and raise the industry's product prices, the sector may become neutral to oil price changes or may even benefit from them. The nature of the expectation formation mechanism also plays a role in determining the effect. In brief, since these industries are quite varied in terms of size and character, they are expected to react differently to macroeconomic changes. Our results manifest these dissimilarities in terms of the delay period, size and sign of the threshold and the changes in the magnitude of the oil return coefficient when oil returns fall below or exceed the threshold values. Based on the symmetry tests reported in Table 17.2, some sectors do show asymmetric responses to oil return changes in the calm period of 1991–1998 with the number increasing considerably in the more volatile 1998–2006 period. We examine only a sample of the oil-consuming industries below.

The results for the Building Construction industry are reported in column 7 of Tables 17.3(A) and 3(B). The threshold value and the coefficients of the oil return effects for the calm and volatile periods are $r = 1.62\%$, $b_4^1 = 0.015$, $b_4^2 = 0.043$ (insignificant) and $r = 1.75\%$, $b_4^1 = -0.014$, $b_4^2 = -0.075$, respectively. The threshold for this industry is rather large in both periods. Moreover, the response of the industry to oil return shocks is symmetric in the first but asymmetric in the second period (Table 17.2). These figures indicate that the industry either benefits from or remains unaffected by oil return increases when oil price markets are relatively steady (calm period), while it suffers when oil prices are highly volatile, especially when oil price increases exceed the threshold. The positive sign of the return on the oil coefficient during the first period may occur because when oil prices are steady for a length of time, building constructors may be better able to pass on the higher cost of fuel to the home buyers. The builders may have also effectively hedged the increase in oil prices.

The situation is quite dissimilar in the 1998–2006 period, when oil prices are higher and highly volatile. In this period, the industry suffers from the oil price increases and its losses only multiply in magnitude from -0.014 (1.4 bp) to -0.075 (7.5 bp) once they pass the threshold. Specifically, during this period, any increase in the price of oil, regardless of its being below or above the threshold, contributes to lower profitability of the Building Construction sector, which in turn leads to a decline in sector return. This scenario is compatible with the commonly expected oil price effects on the oil-consuming sectors. The strong negative effect in this period, especially for oil price increases above the threshold, may be because these increases were generally large and, hence, could not be fully passed on to the consumers. Moreover, home prices had been rising throughout the 1990s and early 2000 and were expected, at least by some observers, to decline in the coming periods. This may have made real estate buyers reluctant to

* Southwest Airlines' experience exemplifies how hedging can be beneficial or detrimental to the profitability of a firm. Southwest hedged against higher fuel prices and purchased long-term contracts to buy most of its fuel at \$51 a barrel through 2009. The value of these hedges soared as oil prices continued to increase throughout 2006 and 2007. However, in 2008, oil prices dramatically decreased to \$40 a barrel. Consequently, Southwest lost \$56.0 million, or 8 cents per share in the fourth quarter of 2008 (www.forbes.com).

[†] Higher oil prices provide more investment opportunities for various sectors. For example, in the oil exploration area, higher prices permit exploration of oil in remote areas/deep sea that were not cost efficient before. Higher oil prices also lead to product innovations and development of substitutes for oil-using products. The auto industry is a perfect example of this phenomenon. As oil prices increase, consumers flock to electric and plug-in-hybrid vehicles. Currently, the auto industry views hybrid cars as the next big profit center. The U.S. is the largest hybrid car market in the world, with sales accounting for 60–70% of global hybrid sales. According to J.D. Power and Associates, hybrid-electric vehicle sales volumes in the country are expected to grow 268% between 2005 and 2012. Presently, there are only 12 hybrid models available in the U.S., which would increase to 52 by 2012 (<http://www.dailymarkets.com/stock/2011/05/24/analyst-interviews-auto-industry-outlook-and-review-4/>).

absorb the price increments due to higher oil prices. Overall, given that oil price increases were high and volatile and they occurred when home prices had been rising for a good length of time, oil price changes were heavily reflected in profitability and stock returns of the building industry. These price increases may have also been larger than expected and not hedged by the industry.

Another example of a major oil-consuming industry is Air Transportation (column 8), for which the threshold and the oil return coefficients for the two periods are, respectively, $r = 1.15\%$, $b_4^1 = -0.025$, $b_4^2 = -0.050$ and $r = 2.31\%$, $b_4^1 = -0.092$, $b_4^2 = -0.088$. The effect of oil on this industry is strong and clear-cut. A few points are notable. First, threshold values are positive and large, indicating that this industry requires a substantial increase in oil prices for a regime change. Second, oil shock effects are universally negative, indicating a loss due to oil price increases, regardless of their being below or above the threshold, or during the calm or volatile period. This result makes sense because fuel is a major cost item for the airline industry. Third, the magnitude of the oil effect is much greater in the volatile period than in the calm period, perhaps because the average price is much higher and the industry may be unable to pass on further increases in the cost of fuel to its clients. Fourth, when the price increases exceed the threshold, the effect on the industry doubles in the low volatility period, although it remains rather unchanged in the high volatility period because the effect is already very high. In other words, given the strong sensitivity of this sector to oil prices in this latter period, the industry shows symmetric effects or equal vulnerability to oil return changes when these returns are below and above the threshold.

For the Chemical, Plastic and Transportation Equipment sectors (columns 9–11) the effects are mostly insignificant when oil returns fall below the threshold, but mostly significant when they exceed the threshold. The direction of the effect for the significant cases is negative for Chemical and Plastic but positive for the Transport Equipment sector. Since these sectors are not oil-intensive, these findings are reasonable. These results may have also received a contribution from the market demand and supply conditions and the relative market power of each industry, which may or may not allow the increased costs of fuel to be passed on to the consumers of these industries' products. Asymmetry of oil shock effects holds for both periods (Table 17.2).

17.3.3.1.4 Depository Institution (DI) Sector

The threshold and the oil return coefficients for the DI sector for the two periods are, respectively, $r = 1.58\%$, $b_4^1 = -0.005$, $b_4^2 = 0.009$ (insignificant) and $r = -0.56\%$, $b_4^1 = -0.013$ and $b_4^2 = -0.015$ (column 12). Unlike the sectors discussed earlier, DIs do not produce or consume significant amounts of oil as a part of their core business. Their profitability is derived from their exposure to loans, deposits and investments in businesses and households affected by oil price changes. Hence, changes in the composition of the loans, investments and deposits and changes in the spreads and the default rates of the loans, due to oil price shocks, largely determine the sensitivity of the DI's premium to the oil return. Thus, the DI sector may benefit or suffer from oil price inflation depending on the mix of the borrowers and depositors.

Our findings indicate that, in both periods, there is an inverse relationship between changes in oil return and those of DI stock return, although in one case the effect is insignificant (calm period when oil returns exceed the threshold). Given that there are many more oil-consuming than oil-producing sectors among the DI customers, these findings are reasonable because increased fuel costs increase the probability of customers' loan default. In effect, DIs mirror the behaviour of their oil-consuming clients. Particular banks with large deposits from oil-producing countries or sectors may potentially benefit from positive effects, but they do not seem to dominate the sample.

17.3.3.2 Summary of Sector Returns

To summarize the effect of oil return changes on different industries, it can be stated that, during the less volatile period in the oil market, R_{oil} played a less important role in determining sector returns. During this period (1991–1998), three sectors, Oil–Gas Extraction, Petroleum Refinery and Air Transportation, show significant R_{oil} effects in both regimes ($j = 1, 2$), while the Plastic and Rubber sector displays no effect in either regime and the rest show mixed results (Table 17.3A). For three sectors, Coal, Building

and DIs, the R_{oil} effect is significant only when oil returns are below the threshold ($j = 1$), and for another three sectors, Chemical, Transportation Equipment and Electric-Gas Services, the effect is prevalent only when the oil return exceeds it ($j = 2$). In contrast, for the more volatile oil price period (1998–2006), the R_{oil} sensitivity results, reported in Table 17.3(B), are highly robust (uniform); R_{oil} sensitivity is significant for seven out of the 10 sectors considered in both regimes ($j = 1, 2$). Of the other three sections, Transport Equipment, and Electric–Gas Services sectors show R_{oil} sensitivity only in regime $j = 1$, and the Chemical sector shows R_{oil} sensitivity only in regime $j = 2$. The results of the symmetry of R_{oil} sensitivity during the volatile oil price period also suggest a marked difference in the R_{oil} effect when oil returns are below and above the threshold ($j = 1, 2$). During this period, the R_{oil} sensitivity between $j = 1$ and $j = 2$ is statistically different for six out of the 10 sectors: Oil–Gas, Building, Chemical, Plastic and Rubber, Transport Equipment and Electric–Gas Services sectors. On the other hand, in the calm period, only two sectors, Transport Equipment and Electric–Gas Services, show asymmetric effects.

Moreover, our findings show that there is salient information value for asset pricing in oil return shocks that is not accounted for in existing models. Two points are notable. First, we find that oil return shocks exceeding the threshold motivate slope shifts for all of the risk factors, not just the oil return factor, for most sectors, and more so for the period of oil price volatility. These shifts are assumed away in single-regime models as the oil price sensitivity of each factor in the model is taken to be identical for all shocks. Overlooking these slope shifts may lead to considerable underestimation (overestimation) of the oil shock impact and erroneous policy recommendations or consumer budget allocations. Second, the oil return is found to be an important asset pricing factor in its own right, as discussed below, which is unaccounted for within the framework of the Fama–French three-factor model. The omission of this factor from the model is likely to produce misleading results and incorrect inferences.

17.3.3.3 Economic Significance

To examine the economic significance of the factors in the model, the effect of one standard deviation change in each factor on the return of each sector is calculated by multiplying the coefficient estimate for the factor by its respective standard deviation, reported in Table 17.1. The results are displayed in Tables 17.4(A) and 4(B). According to these results, although the economic significance of the oil return variable falls short of those for the Fama–French factors in both calm and volatile oil price periods, there are some cases in which the economic effect of this variable is large enough to exceed that of the growth (HML) and/or size (SMB) variables. The economic magnitudes of the oil return effects are found to be much larger in the more recent volatile oil price period (e.g. SIC 12 and SIC 45) and these effects more frequently exceed those of the Fama–French growth and size variables (e.g. SIC 13 and SIC 29).

17.3.4 The FIGARCH Volatility Equation

A main contribution of this study is to model the persistence of shocks to the volatility of sector returns using the FIGARCH specification. We study the nature of shock persistence due to oil return changes to determine whether the effects are subject to decay and, if so, whether they follow an exponential (geometric) or a slower hyperbolic rate of decline. The parameters in the FIGARCH (1, δ , 1) volatility equation, ω_0^j , β_0^j , ϕ_0^j and δ^j ($j = 1, 2$), are presented in Tables 17.3(A) and 3(B). The lack of significance of the constant term, the time-invariant component of the volatility equation (ω_0^j), in both sub-periods, indicates that sectoral return volatility is entirely driven by past and current shocks; it does not contain a homoskedastic component.* This stands in sharp contrast to most extant studies that assume return volatility is invariant.

The parameter β_0^j for the autoregressive term σ_{t-1}^2 in the volatility equation is the indicator of short-term volatility decay; the greater the value of β_0^j , the greater the shock persistence (the slower the decay) will be. Values of this parameter less than unity indicate the absence of explosive volatility behaviour, as

* This finding is similar to those obtained by Baillie *et al.* (2002) in modelling U.S. inflation.

TABLE 17.4A Economic Significance of Double-Threshold FIGARCH(1, δ , 1) Results January 25, 1991 to December 10, 1998 Sub-Period.

1	2	3	4	5	6	7	8	9	10	11	12
Coefficient	Variable	Coal (SIC12)	Electric-Gas Services (SIC49)	Oil-Gas Extraction (SIC13)	Petroleum Refinery (SIC29)	Building (SIC15)	Air Transportation (SIC45)	Chemical (SIC28)	Plastic and Rubber (SIC30)	Transport Equipment (SIC37)	Depository Inst. (SIC60)
b_1^2	RM	0.754	0.573	0.805	0.658	1.056	1.144	0.962	0.685	0.912	0.684
b_1^1	RM	0.857	0.553	0.804	0.704	1.043	1.090	0.856	0.660	0.859	0.632
b_2^2	SMB		0.175	0.287	0.187	0.757	0.675	0.519	0.386	0.410	0.306
b_2^1	SMB	0.330	0.125	0.432	0.150	0.533	0.422	0.408	0.363	0.338	0.275
b_3^2	HML	0.209	0.167	0.130	0.272	0.283	0.322	-0.008	0.112	0.241	0.246
b_3^1	HML	0.206	0.212	0.201	0.239	0.364	0.181	0.018	0.156	0.185	0.245
b_4^2	R_{oil}		0.007	0.166	0.094	0.101	-0.117	-0.016	0.012	-0.012	0.021
b_4^1	R_{oil}	0.035	-0.023	0.148	0.082		-0.059	-0.016		0.019	

TABLE 17.4B Economic Significance for Double-Threshold FIGARCH(1, δ , 1) Results December 11, 1998 to December 29, 2006 Sub-Period.

1	2	3	4	5	6	7	8	9	10	11	12
Coefficient	Variable	Coal (SIC12)	Electric-Gas Services (SIC49)	Oil-Gas Extraction (SIC13)	Petroleum Refinery (SIC29)	Building (SIC15)	Air Transportation (SIC45)	Chemical (SIC28)	Plastic and Rubber (SIC30)	Transport Equipment (SIC37)	Depository Inst. (SIC60)
b_1^2	RM	0.708	0.623	0.854	0.863	0.917	1.288	0.751	0.693	0.766	0.427
b_1^1	RM	0.824	0.713	0.757	0.815	1.003	0.979	0.894	0.599	0.837	0.393
b_2^2	SMB	0.253	0.109	0.298	0.115	0.346	0.343	0.429	0.208	0.206	0.136
b_2^1	SMB	0.358	0.136	0.296	0.197	0.376	0.306	0.459	0.252	0.267	0.142
b_3^2	HML	0.393	0.352	0.494	0.504	0.178	0.485	0.067	0.216	0.221	0.183
b_3^1	HML	0.452	0.395	0.436	0.439	0.410	0.323	0.106	0.228	0.360	0.139
b_4^2	R_{oil}	0.241		0.464	0.248	-0.176	-0.206	-0.033	-0.068		-0.035
b_4^1	R_{oil}	0.195	0.026	0.506	0.312	-0.033	-0.216		0.016	-0.040	-0.030

well as the failure of the constant variance models in accurately describing the sectoral stock indexes.[†] Oil shocks affect both the real and financial sectors of the economy, altering the way investors perceive risk and form expectations in response to the arrival of oil price news. The coefficient ϕ_0^j , expected to lie between one and zero ($1 \geq \phi_0^j \geq 0$), describes the moving-average (MA) component of the short-term volatility dynamics and provides an approximation to the level of short-term shock persistence as its value affects the degree of dependence of the current volatility on lagged squared errors. Although within a FIGARCH model $\phi_0^j = 1$ suggests an IGARCH process, Baillie *et al.* (1996) point out that the interpretation of the decay process is substantially different for these two models. Specifically, while IGARCH imposes complete persistence in its impulse response weights, FIGARCH implies a long-memory behaviour (a slow rate of decay) as determined by the magnitude of δ^j , producing a mean-reverting pattern. In the low volatility period, all ϕ_0^j values, with the exception of the Building and Air Transportation sectors ($j = 2$, $\phi_0^j = 1.045$ and 1.285 , respectively), are less than unity, satisfying the regularity conditions. In four cases (three in regime 1 and one in regime 2), ϕ_0^j is insignificant, indicating that the memory does not include a moving-average component.[‡] The results are rather similar between low and high volatility periods.

The finding of regime shifts in the volatility equation at threshold oil return, evidenced by the dissimilarity of the volatility equation in the two regimes, denotes the existence of oil-driven structural changes in the parameters of the volatility equation, rendering the models overlooking this feature suspect. It follows that the short-term volatility structure of sector returns is sensitive to changes in the oil industry, and the estimates of oil shock effects based on simpler models are likely to be distorted. Previous research (e.g. Hamilton and Susmel (1994) and Gray (1996)) has documented that overlooking structural changes in the GARCH parameters creates spuriously high persistence in conditional volatility, as well as poor forecasting performance. There are also economic justifications for allowing GARCH parameters to vary across regimes as some industries, e.g. the financial sector, are able to adjust to shocks quickly, while others, e.g. Coal Mining, take much longer to absorb them because of the nature of their production processes. Under a single-regime specification, the parameters of the conditional covariance matrix are assumed to stay constant over time. The threshold approach explicitly recognizes the possibility of shifts in these parameters.

Parameter δ^j is a measure of long-term memory decay. As it captures the long-range dependence in volatility, larger values of this parameter indicate a slower pace of decay, or a higher level of persistence, in the long run. In particular, when this parameter takes a value of unity, the model is described as an IGARCH process according to which shocks persist permanently, while a zero value of this parameter reduces the model to a GARCH process with a more rapid exponential pace of decay. The presence of long memory (hyperbolic decay) in the conditional variance of sector returns is strongly supported by the magnitude and significance of the fractional difference coefficients δ^1 and δ^2 . The decay parameter δ^j is significant and in an acceptable range for all sectors in both periods with one exception (Petroleum Refinery in the first period in regime $j = 1$). Thus, the FIGARCH coefficients simultaneously present the long-term dynamics captured by significant decay parameters, and the short-term dynamics represented by significant autoregressive and moving-average parameters (β and ϕ). Large and significant values for $0 < \delta < 1$ imply strong support for the hyperbolic decay and condition of high persistence, as opposed to the conventional exponential decay associated with the GARCH (1, 1), or indefinite persistence with the IGARCH (1, 1) processes. Additional support for hyperbolic decay is provided by the

* Only the GARCH parameter for the Air Transportation sector exceeds the unity value in regime $j = 2$ during the calm period (β_0^2), suggesting that its conditional volatility is non-stationary in the short run. However, this path is truncated at periodic intervals by the switch to the stationary regime as exhibited by the β_0^j value being less than unity in regime $j = 1$. A chi-square test of equality of the GARCH parameters ($H_0: \beta_0^1 = \beta_0^2$) is rejected at the 1% level.

† Other specifications, including GARCH (1, 1) and FIGARCH (1, δ , 0), were also tested. The former model fails to converge in a number of cases. The latter produces results similar to those obtained above.

‡ In such a situation, a FIGARCH (1, δ , 0) process may better model the return volatility. Even though ϕ_0^j is found to be insignificant, the fractional difference parameter, δ , still properly captures the long-memory decay process.

rejection of hypotheses H_{11} (GARCH model: $\delta^1 = \delta^2 = 0$) and H_{12} (IGARCH model: $\delta^1 = \delta^2 = 1$), the results for which are reported in Table 17.2. These results are consistent with, and reinforce, the finding of long memory for squared index returns and individual stock returns of Lobato and Savin (1998).

Further analysis of the long-range volatility dependence parameters δ^1 and δ^2 reveals that the pattern of long-term decay in response to volatility shocks is regime-dependent because the coefficient δ^1 pertaining to oil return changes below the threshold ($j = 1$) is larger in magnitude than its counterpart δ^2 when oil returns exceed the threshold ($j = 2$), both in the first and second periods. This indicates that when oil returns are below the threshold the decay is slower and shock effects are more persistent, regardless of the degree of oil price volatility. The dissimilarity of the decay process under the two regimes is further strengthened by the rejection of the asymmetry tests (H_{25} : $\delta^1 = \delta^2$) discussed earlier. The findings of regime shifts and the dissimilarity of parameter values under calm and volatile oil price periods also support the view that oil price shocks influence investor perception of risk in the short run as well as bringing about structural changes in risk determination processes and thereby exerting profound influences on the way investors perceive risk and form expectations in the long run.

17.4 Conclusions

This study is the first to provide a detailed assessment of the effects of oil return changes on the equity return distributions of 10 major industrial and financial sectors in the U.S. within the framework of a Double-Threshold FIGARCH (1, δ , 1) model. This framework allows us to delineate the magnitude and persistence of oil shock effects at the sector level, which can help in the assessment of sectoral diversification benefits by investors, hedging decisions by firm managers and the formulation of regulatory actions by government. The issue of sectoral effects deserves special attention, in particular because market and country factors have become more and more correlated over time, due to technological advancement, deregulation and globalization, and thus less helpful in diversifying portfolio risk, while sector-specific risk continues to play an important role in diversification, investment and risk management decisions.

This study addresses the long-memory property in the return volatility of various sectors in the U.S. in response to oil shocks within a generalized asset pricing model including the Fama–French factors together with return on oil as explanatory variables. We find strong support for the proposition that the return on oil (R_{oil}) improves upon the basic Fama–French factor model in describing asset returns. Unlike extant studies, we introduce and identify an oil return threshold for each of the sectors considered, beyond which the impacts of the explanatory variables in the model change in magnitude and, sometimes, even in terms of the direction of the effect. We also find that GARCH, IGARCH and Fama–French models are inadequate for describing sectoral returns and accounting for long-term persistence and asymmetric responses to oil shocks.

The threshold effects of oil return changes are sector-specific, as well as period-specific. In particular, during the period of volatile oil markets with an upward trend in oil prices, the threshold values are predominantly positive, whereas during the less volatile and downward trend in the price of oil the threshold values have mixed signs. Moreover, we find that during the less volatile period in the oil market, oil prices played a less significant role in determining sector returns in comparison with the period when oil prices were rising and more volatile. Oil prices also played a less important role when they were located below the threshold rather than above it.

Hypotheses regarding the symmetry of coefficients for above and below threshold oil return levels are largely rejected for the factors included as explanatory variables in sector return models. For the mean equation coefficients, threshold asymmetry is evident for all risk factors in the model, market, size, value and R_{oil} , although less frequently significant for the latter variable. These findings support the need for separating volatile and less volatile periods and the superiority of a threshold specification for modelling sectoral returns over single-regime specifications. We also extend the literature by analysing the conditional volatility of returns. The results reveal the concurrent prevalence of long-term volatility

dynamics (δ) and short-term volatility dynamics (β). The finding of highly significant δ values between $0 < \delta < 1$ supports the stationarity of the volatility process and establishes the validity of the hyperbolic decay (high shock persistence), as opposed to the conventionally used GARCH (1, 1) with exponential (rapid) decay, and the IGARCH (1, 1) process indicating indefinite persistence. This suggests that the appropriate methodology to model sector returns is FIGARCH (1, δ , 1). Finally, the results show that the regime shifts and long-term structural changes caused by oil price shocks have a profound influence on the way investors perceive long-term risk and form their expectations.

Acknowledgements

This paper was presented at the 7th Annual International Conference on Business at the Athens Institute for Education and Research, Athens, Greece. The authors would like to thank the participants at the conference for comments. All errors are ours.

References

- Awerbuch, S. and Sauter, R., Exploiting the oil–GDP effect to support renewable deployment. SPRU Working Paper No. 129, October 2004, 2005.
- Baillie, R.T., Bollerslev, T. and Mikkelsen, H.O., Fractionally integrated generalized autoregressive conditional heteroskedasticity. *J. Econometr.*, 1996, **74**, 3–30.
- Baillie, R.T., Han, Y.W. and Kwon, T.-G., Further long memory properties of inflationary shocks. *Southern Econ. J.*, 2002, **68**(3), 496–510.
- Brooks, C., A double-threshold GARCH model for the French franc/Deutschmark exchange rate. *J. Forecast.*, 2001, **20**, 135–143.
- Chen, N., Roll, R. and Ross, S.A., Economic forces and the stock market. *J. Business*, 1986, **59**, 383–403.
- Choi, K. and Hammoudeh, H., Volatility behavior of oil, industrial, commodity and stock markets in a regime-switching environment. *Energy Policy*, 2010, **38**, 4388–4399.
- Davis, S.J. and Haltiwanger, J., Sectoral job creation and destruction response to oil price changes. *J. Monet. Econ.*, 2001, **48**, 465–512.
- Dickey, D. and Fuller, W., Distribution of the estimators for autoregressive time series with a unit root. *J. Am. Statist. Assoc.*, 1979, **74**, 427–431.
- Dickey, D. and Fuller, W., Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, 1981, **49**, 1057–1072.
- Driesprong, G., Jacobson, B. and Benjamin, M., Striking oil: Another puzzle? *J. Financ. Econ.*, 2008, **89**, 307–327.
- Elyasiani, E., Mansur, I. and Odusami, B., Oil price shocks and industry stock return. *Energy Econ.*, 2011, **33**(5), 966–974.
- Faff, R. and Brailsford, T., Oil price and the Australian stock market. *J. Energy Finance Devel.*, 1999, **4**, 69–87.
- Fama, E. and French, K., The cross-section of expected stock returns. *J. Finance*, 1992, **47**, 427–465.
- Fama, E. and French, K., Industry costs of equity. *J. Financ. Econ.*, 1997, **43**, 153–193.
- Gray, S.F., Modeling the conditional distribution of interest rates as a regime switching process. *J. Financ. Econ.*, 1996, **42**, 27–62.
- Greene, D.L. and Tishchishyna, N.I., *Costs of Oil Dependence: A 2000 Update*, 2000 (Oak Ridge National Laboratory: Oak Ridge, Tenn.).
- Guo, H. and Kliesen, K.L., Oil price volatility and U.S. macroeconomic activity. *Fed. Res. Bank St. Louis Rev.*, 2005, **87**(6), 669–683.
- Hamilton, J.D., What is an oil shock? *J. Econometr.*, 2003, **113**, 363–398.
- Hamilton, J.D., Oil and the macroeconomy. In *The New Palgrave Dictionary of Economics*, 2nd ed., edited by S.N. Durlauf and L.E. Blume, 2008 (Houndmills: U.K./ Palgrave Macmillan: New York).

- Hamilton, J.D. and Susmel, R., Autoregressive conditional heteroskedasticity and change in regime. *J. Econometr.*, 1994, **64**, 307–333.
- Hammoudeh, S., Yuan, Y., Chiang, T. and Nandha, M., Symmetric and asymmetric US sector return volatilities in presence of oil, financial and economic risks. *Energy Policy*, 2010, **38**, 3922–3932.
- Hooker, M.A., Are oil shocks inflationary? Asymmetric and nonlinear specifications versus changes in regime. *J. Money, Credit Bank.*, 2002, **34**, 540–561.
- Huang, B., Hwang, M.J. and Peng, H., The asymmetry of the impact of oil price shocks on economic activities: An application of the multivariate threshold model. *Energy Econ.*, 2005, **27**, 455–476.
- Huang, R.D., Masulis, R.W. and Stoll, H.R., Energy shocks and financial markets. *J. Fut. Mkt*, 1996, **16**, 1–27.
- Jarque, C.M. and Bera, A.K., A test for normality of observations and regression residuals. *Int. Statist. Rev.*, 1987, **55**, 163–172.
- Jones, C.M. and Kaul, G., Oil and the stock markets. *J. Finance*, 1996, **51**, 463–491.
- Keane, M.P. and Prasad, E.S., The employment and wage effects of oil price changes: A sectoral analysis. *Rev. Econ. Statist.*, 1996, **78**, 389–400.
- Kilian, L., Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *Rev. Econ. Statist.*, 2008a, **90**, 216–240.
- Kilian, L., The economic effects of energy price shocks. *J. Econ. Liter.*, 2008b, **46**, 871–1009.
- Kilian, L., Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *Am. Econ. Rev.*, 2009, **99**, 1053–1069.
- Lee, K.S. and Ni, S., On the dynamic effects of oil price shocks: A study using industry level data. *J. Monet. Econ.*, 2002, **49**, 823–852.
- Li, C.W. and Li, W.K., On a double threshold autoregressive heteroscedastic time series model. *J. Appl. Econometr.*, 1996, **11**, 253–274.
- Ljung, G.M. and Box, G.E.P., On a measure of a lack of fit in time series models. *Biometrika*, 1978, **65**, 297–303.
- Lobato, I. and Savin, N.E., Real and spurious long-memory properties of stock market data. *J. Business Econ. Statist.*, 1998, **16**, 261–283.
- Meyer, J. and von Cramon-Taubadel, S., Asymmetric price transmission: A survey. *J. Agric. Econ.*, 2004, **55**, 581–611.
- Mork, K.A., Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *J. Polit. Econ.*, 1989, **97**, 740–744.
- Moskowitz, T. and Grinblatt, M., Do industries explain momentum? *J. Finance*, 1999, **54**, 1249–1290.
- Perron, P., Trends and random walks in macroeconomic time series: Further evidence from a new approach. *J. Econ. Dynam. Control*, 1988, **12**, 297–332.
- Perron, P., The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 1989, **57**, 1361–1401.
- Phillips, P.C.B. and Perron, P., Testing for a unit root in time series regression. *Biometrika*, 1988, **75**, 335–346.
- Poterba, J. and Summers, L., The persistence of volatility and stock market fluctuations. *Am. Econ. Rev.*, 1986, **76**, 1142–1151.
- Sadorsky, P., Oil price shocks and stock market activity. *Energy Econ.*, 1999, **21**, 449–469.
- Sadorsky, P., Risk factors in stock returns of Canadian oil and gas companies. *Energy Econ.*, 2001, **23**, 17–28.
- Tsay, R.E., Testing and modeling threshold autoregressive process. *J. Am. Statist. Assoc.*, 1989, **84**, 231–240.
- Tsay, R.E., Testing and modeling multivariate threshold models. *J. Am. Statist. Assoc.*, 1998, **93**, 1188–1202.
- Yang, C.W., Hwang, M.J. and Huang, B.N., An analysis of factors affecting price volatility of the US oil market. *Energy Econ.*, 2002, **24**, 107–119.
- Zivot, E. and Andrews, D.W.K., Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *J. Business Econ. Statist.*, 1992, **10**, 251–270.

Appendix A: Threshold and Delay Parameter Values

In a threshold model, a smaller threshold value for a sector means that a comparatively smaller increase in oil return is sufficient to trigger a regime shift for that sector. Threshold values can be positive or negative. In the latter case a regime change occurs when the oil return is falling. This is possible because the slowing of oil price declines can alter the expectations about future changes in oil prices, or even the expectation generation mechanism itself. For example, if expectations are extrapolative and subject to a momentum effect, slowing of oil price reductions may lead to expectations of further slowing and even a reversal in the oil price trend towards a continuum of price increases. In this scenario, a negative threshold is reasonable because a slowing in the oil price declines can be sufficient to alarm some investors about future oil price increases, prompting a selloff (or additional purchases) of the stocks in the sectors being harmed (bolstered) by the change in the oil price patterns.

Threshold values may be larger or smaller during highly volatile periods (1998–2006) compared with calm periods (1991–1998). If high oil price volatility leads economic agents to consider the oil price changes as mostly transitory, it will take more of a change in oil prices to trigger regime shifts and the effect will also decay more quickly. In this case the threshold value (r) in a volatile period may be larger than that for a calm period. Conversely, if several months of steady oil price levels in a calm period create a perception that steady prices are the norm and any changes are temporary, then the threshold may be larger during the calm period. Smaller threshold values may occur when expectations about a price increase are firmer (more certain), price changes are widely considered to be permanent, and/or agents are less tolerant of these changes. In general, a relatively smaller or larger threshold may be manifested during volatile or calm periods, depending on the prevailing expectation formation mechanism and characteristics of the industry. Delays in regime shifts (d) may also be shorter or longer depending on the same factors, as well as the speed of information transmission in the financial markets, and the composition of market players between individual and institutional, and between domestic and foreign investors. Change in the factor slopes (regime shifts) may be upward or downward. That is, the slopes b_i^1 , in effect for oil return variations above the threshold, can be smaller or larger than b_i^2 , for changes below the threshold. When b_i^1 is smaller in value than b_i^2 , sectoral sensitivity to the factors increases when we move above the threshold. On the contrary, if b_i^1 exceeds b_i^2 , once we move above the threshold, there will be a downward movement in factor sensitivities. Fama and French (1997) have provided some theoretical explanations for these slope shifts.

18

Short-Term and Long-Term Dependencies of the S&P 500 Index and Commodity Prices

Michael Graham

Jarno Kiviaho

Jussi Nikkinen

18.1	Introduction	388
18.2	Brief Literature Review	389
18.3	Measuring Stock and Commodities Market Dependencies Using Wavelet Squared Coherency	390
18.4	Data	391
	Commodity and Stock Indexes • Data Description	
18.5	Empirical Analysis of Comovements	393
18.6	Conclusion	398
	Acknowledgement	399
	References	399

We utilize wavelet coherency methodology with simulated confidence bounds to examine the short-term and long-term dependencies of the returns for S&P 500 and the S&P GSCI® commodity index. Our results indicate no evidence of comovement between S&P 500 total return and the S&P GSCI® commodity index total return in the short term, thereby suggesting diversification gains for equity investors. Importantly, this finding encompasses the onset of the current financial crisis. However, long-term diversification benefits, particularly after the onset of the recent financial crisis, are limited. We find, moreover, no consistent evidence of comovements between S&P 500 and 10 individual sub-indexes of the S&P GSCI® commodity index. Of particular importance, we report weak comovement of returns between S&P 500 and S&P GSCI® Precious Metals total return and S&P 500 and S&P GSCI® Softs at all frequencies, implying significant diversification gains both for short-term and long-term investors.

Keywords: Commodity markets; Comovement; Applied finance; Correlation modelling

JEL Classification: C4, E3, E32, F3, F30, G1, G15

18.1 Introduction

The volatility in global markets emanating from the recent financial crisis has presented real and momentous challenges for portfolio management as cross-market dependencies have increased.* In the last two decades alone, financial markets have endured a catastrophic collapse in confidence resulting from the Mexican, Asian and Russian crises. The most recent financial crisis, which started in late 2007, has been different from these earlier crises given its global and pervasive impact across different markets. Many studies have shown that the various asset markets have all been significantly affected by the recent events in unforeseen ways that are not mirrored by the experiences of past crises.

In this paper, we investigate whether commodities could potentially provide some diversification benefits for the stock market investor during periods of financial turmoil. We do this by examining the short-term and long-term comovement of the returns series for S&P 500 and the S&P GSCI® commodity index using wavelet coherency analysis. We also examine dependencies between the S&P 500 and 10 sub-indexes of the S&P GSCI® commodity index.†

The seminal portfolio theory of Nobel Laureate Professor Harry Markowitz (1952) implicitly posits that uniquely including assets with less than perfect correlations in a portfolio reduces its inherent diversifiable risk. Commodities as an asset class has shaken off its reputation as a comparatively unknown asset class to become a new unique asset class (see Gorton and Rouwenhorst (2006a,b) for more on this), with the potential to reduce systematic risk in a portfolio. Currently, many academics (see, e.g. Billingsley and Chance (1996), Jensen *et al.* (2002) and Gorton and Rouwenhorst (2006a,b)) and market participants consider commodities to be an important element of being diversified‡ given that they show low correlations (Gorton and Rouwenhorst 2006a,b, Ibbotson Associates 2006) with traditional assets (e.g. equities and bonds) that are robust to extreme events. This is because factors that influence commodity prices (e.g. weather and the geopolitical situation, supply constraints in physical production and event risk) differ from those that affect the value of traditional assets (Geman 2005). Significantly, assets tied to the S&P GSCI index as of 31 December 2010 amounted to a staggering \$US100 billion.§ Following this intense interest in commodities, the financial press have also keenly reported on this market and published many articles discussing the role of commodities in diversified portfolios.¶

The recent global financial crisis has raised concerns, particularly among some market observers, that the diversification argument in favour of commodities investment is weak. For example, in the fourth quarter of 2008, financial markets experienced simultaneous price declines in every financial asset, with the exception of the U.S. Treasuries. Tang and Xiong (2010) and Silvennoinen and Thorp (2010) suggest that the growing prominence of commodities index investors increases comovement between commodity markets and traditional asset markets. Kyle and Xiong (2001) also posit that portfolio rebalancing by commodity index investors can spill over price volatility from other markets to commodity markets, potentially exposing commodities to time-scale dependencies with traditional asset markets. Given the

* Financial crisis interferes with the resource allocation process that moves funds to agents with the most productive investment opportunities in the financial system.

† This paper deals with dependence in time of crisis. Other studies have concentrated on extreme comovement in international financial markets. See Longin and Solnik (2001), Ang and Chen (2002), Patten (2004), Garcia and Tsafack (2009), Adel and Salma (2012) and references therein for a discussion of this literature.

‡ The value of commodities is best achieved as part of a diversified portfolio and not as a stand-alone investment. Historically, commodities have shown inferior returns and greater volatility when compared with other assets (Erb and Harvey 2006). However, when combined with conventional investment assets (e.g. stocks and bonds), they can enhance portfolio performance. See also, for instance, the article of Mark Huamani on investing in commodities and diversification: http://www.jpmorgan.com/tss/General/Investing_in_Commodities/1159337364479 (accessed 1 September 2011).

§ <http://www.bloomberg.com/news/2011-01-07/s-p-gsci-says-commodity-assets-against-index-rose-as-high-as-100-billion.html> (accessed 1 September 2011).

¶ See, for instance, <http://www.ft.com/intl/cms/s/1/23f4dee4-4803-11db-a42e-0000779e2340.html#axzz1YDqkOknx> (an article by John Authers, *Financial Times*, September 22, 2006) (accessed 8 November 2010).

importance and practical implications of asset dependencies in asset allocation decisions and risk management, we use an enhanced methodology, the three-dimensional analysis of wavelet coherency, to investigate the comovement of the returns series for S&P 500 and the S&P GSCI® commodity index for a 10-year period encompassing the height of the current global financial crisis. Low comovement during the crisis would imply that investing in commodities as part of a diversified equity portfolio can reduce portfolio volatility during periods of financial crisis.

Our empirical investigations add to the body of knowledge on stock market and commodities market dependencies and offer portfolio managers novel evidence of the observed stock and commodity market dependencies. The current literature intimates that the commodity market and equity market dependencies vary over time (Chance 1994). However, studies examining the usefulness of commodities as part of a portfolio comprising traditional financial assets typically make no distinction between short-term and long-term investors in their analyses (see, e.g. Jensen *et al.* (2000, 2002), Gorton and Rouwenhorst (2006b), Ibbotson Associates (2006) and Yamori (2011)). Thus the observed time variation of asset dependencies lacks an explicit interpretation from the point of view of investors in different time horizons. This difference is important because short-term and long-term investors would be interested in different fluctuations in the comovements of the commodities and equity markets (Candelon *et al.* 2008).^{*} We investigate this difference and explore the potential time-varying behaviour of stock and commodities market dependencies in a unified empirical framework.

The wavelet coherency methodology employed in this study offers a refinement in terms of analysis as both time and frequency domains are simultaneously taken into consideration to measure market dependencies in an integrated framework. Accordingly, both transient and long-term associations between stock and commodity markets can be analysed concurrently. The output of this empirical tool enables researchers and analysts to zero in on regions where equity market and commodity market dependencies are higher for both short-term and long-term investors. In such regions, the benefits of portfolio diversification, in terms of risk management, are lower.

The evidence presented implies that combining an equity portfolio with commodities can provide diversification benefits in the short term. The dependencies between the S&P 500 and S&P GSCI® commodity indexes as well as between S&P 500 and 10 sub-sets of the S&P GSCI® commodity index all show weak fluctuations at the highest frequencies. Moreover, this result is shown to be robust during the current financial crisis. In the long term, however, our results show that there are relatively reduced gains from combining an equity portfolio with commodities as the onset of the current financial crisis ushered in a period of increased dependencies between stock index returns and commodity index returns. These results are, in general, confirmed when we examine comovements between the equity portfolio, S&P 500, and 10 sub-indexes of the commodity index, and the S&P GSCI® commodity index, but with notable exceptions. For example, the evidence presented also shows enhanced diversification benefits by combining equity (S&P 500 total return) and the S&P GSCI® Precious Metal total return in the long term as well as in the short term.

The remainder of the paper is organized as follows. Section 18.2 briefly reviews the relevant literature on commodity and portfolio diversification. Section 18.3 presents the wavelet methodology used to analyse comovements of stock and commodity market returns. The data employed in the study and our empirical findings are discussed in Sections 18.4 and 18.5, respectively, and Section 18.6 concludes.

18.2 Brief Literature Review

The identification of dependencies between financial assets is a key ingredient in risk assessment and portfolio management. Modern portfolio theory demonstrates that the value of diversification is bounded by the systematic risk of the market as a whole. Gorton and Rouwenhorst (2006a, p. 20),

^{*} Utilizing wavelet coherency methodology is not novel, however. The methodology has been applied to analyse financial time-series (see, e.g. Rua and Nunes (2009), Nikkinen *et al.* (2011), Graham and Nikkinen 2011) and Graham *et al.* (2012)).

however, point out that commodities have *some power at diversifying the systematic component of risk*. Additionally, this asset class is strongly correlated to the inflation rate and changes in the rate of inflation, which provides a hedge against rising inflation (see, for example, Anson (1999)). It is therefore not surprising that vast amounts of money have been invested in commodity linked indexes by hedge funds and pension funds (Akey 2005).

Numerous academic studies have investigated the benefits of commodities for portfolio diversification and as a source of investment returns. For instance, Lummer and Siegel (1993), Ankrum and Hensel (1993) and Kaplan and Lummer (1998) find that allocating resources to this asset class offers excellent diversification benefits for equity and bonds in a mean–variance static asset allocation framework. Fortenberry and Hauser (1990) and Conover *et al.* (2010) also note the risk reduction benefit of including commodities in a portfolio of traditional assets without sacrificing returns. Furthermore, Satyanarayan and Varangis (1996), Abanomey and Mathur (1999) and Garreth and Taylor (2001) find enhancements in portfolio returns for a given level of risk (an upward shift in the efficient frontier) if an allocation of commodities is included in an international stock portfolio. Similarly, Greer (1994) finds that the effect of equal weights invested in S&P 500 and commodities produces superior portfolio returns with a lower standard deviation relative to investing only in the S&P 500. Anson (1999), Jensen *et al.* (2000, 2002), Ibbotson Associates (2006), Gorton and Rouwenhorst (2006b), Idzortek (2007) and Nguyen and Sercu (2010) all note the diversification advantages, in terms of portfolio performance, of commodity indexes. Additionally Gorton *et al.* (2005) and Belousova and Dorfleitner (2012) generalize the U.S.-based conclusion to Japanese and European investors.

Erb and Harvey (2006), exploring the tactical and strategic opportunities that the inclusion of commodities in traditional portfolios present, however, cast doubt on the role of commodities in long-term asset allocation. Furthermore, Cheung and Miu (2010) show that the diversification benefit of including commodities in a portfolio of traditional assets is not convincing. Similarly, Daskalaki and Skiadopoulos (2011) find that optimal portfolios that include only the traditional assets show superior performance, thus challenging earlier findings advocating the benefits of commodities in investors' portfolios. Given the mixed results, an examination of equity and commodity market comovement utilizing an integrated framework offers a fresh understanding of the nature of the dependency between the two asset classes and its implications for asset allocation.

18.3 Measuring Stock and Commodities Market Dependencies using Wavelet Squared Coherency

Wavelet coherency analysis is used to examine dependencies between the stock market and the commodity market returns. We briefly review the concept in this section.

Wavelet analysis is a mathematical tool for extracting scale/frequency information from data (see, e.g. Daubechies (1992) or Mallat (1998) for a thorough exposition). In contrast to the other popular frequency analysis method, e.g. Fourier analysis, wavelets preserve the time-localized information of the data and are therefore ideal for analysing non-stationary time-series. Wavelet transforms can be categorized into the discrete wavelet transform and its continuous version. The output from the continuous transform is more redundant than its discrete counterpart and thus more readily applicable for feature extraction purposes. Our presentation closely follows Torrence and Compo (1998) and Grinsted *et al.* (2004).

Inputs to wavelet analysis are a time-series x_n , $n = 0, \dots, N$, with equal time steps δt , and a wavelet function $\psi_\sigma(\eta)$, where η is a dimensionless time parameter. The wavelet function must fulfil the conditions of having a zero mean and being localized both in frequency and time. We utilize the popular wavelet function for feature extraction purposes, the Morlet wavelet, as the basis function for wavelet transform. This function enables the identification and isolation of signals, as it affords a balance between localization of time and frequency. In its simple form, the Morlet wavelet is represented as

$$\psi_n(\eta) = \pi^{1/4} e^{i\omega_0 \eta} e^{-0.5\eta^2}, \quad (18.1)$$

with the dimensionless frequency parameter $\omega_0 = 6$ providing a good balance between localization of time and frequency. The wavelet transform has the ability to estimate the spectral characteristics of signals as a function of time and can provide both the time-varying power spectrum and the phase spectrum needed for the calculation of coherence. The continuous wavelet transform, $W_n(s)$, of a discrete sequence x_n at time n and scale s with the wavelet $\psi_0(\eta)$ of choice is defined as the convolution

$$W_n(s) = \frac{1}{\sqrt{s}} \sum_{n'=1}^N x_{n'} \psi_0^* \left[\frac{(n' - n)\delta t}{s} \right], \quad (18.2)$$

where $*$ denotes the complex conjugate. The resulting (complex) wavelet coefficients can then be used to describe certain features of the time-series. By taking the absolute value of the coefficients, the results can be interpreted as the amplitude or *power* of the series: $|W_n(s)|$. The wavelet power spectrum or variance can be obtained by squaring the power: $|W_n(s)|^2$. Since we are interested in the comovements of two time-series (say x_n and y_n), we calculate the cross-wavelet transform: $W_n^{XY}(s) = W_n^X(s)W_n^Y(s)^*$, where W_n^X and W_n^Y are wavelet transforms of x_n and y_n , respectively. The cross-wavelet power $|W_n^{XY}(s)|$ can be interpreted as depicting the local covariance between the two time-series at each scale and frequency. Thus it indicates areas in the scale–time space where the two time-series share high common power.

Finally, to obtain a practical measure of the dependencies of two time-series, we calculate the wavelet squared coherency. This is done by squaring the cross-wavelet power and dividing the result by the product of the two wavelet spectra:

$$R_n^2(s) = \frac{|S(s^{-1}W_n^{XY}(s))|^2}{S(s^{-1}|W_n^X(s)|^2) \cdot S(s^{-1}|W_n^Y(s)|^2)} \quad (18.3)$$

where S is a smoothing operator. Since the resulting squared coherency is normalized to take values between 0 and 1, it is reminiscent of the classic correlation coefficient.

To assess the statistical significance level of the wavelet coherence, we create a background spectrum by simulating a large number of white noise time-series pairs. The mean background power spectrum is assumed to represent the mean power spectrum of the actual time-series. High power in the actual spectrum compared with the background spectrum can then be assumed to be “a true feature with a certain percent confidence” (Torrence and Compo 1998). We use the 95% confidence level as in the original paper.

18.4 Data

18.4.1 Commodity and Stock Indexes

The commodity returns series utilized in this study is the Standard and Poor Goldman Sachs Commodity Index (S&P GSCI®) of commodity sector returns.* S&P GSCI® is a tradable index and a benchmark for investment in the commodity market. It measures the performance in the physical commodities market. The construction of the S&P GSCI® is on a world production-weighted basis, where the quantity of each commodity in the index is determined by the average quantity of production in the last five years

* We acknowledge the existence of alternatives, e.g. public futures funds, commodity trading advisors, and private commodity pool operators, through which investors can invest in commodities (Edwards and Liew 1999). However, using the S&P GSCI® is advantageous because it allows for a clean examination without the contamination of superior/inferior trading techniques of other active public futures funds. Our choice of index also avoids biases associated with databases that report futures returns. Akey (2005) provides an excellent review of available passive commodity indexes.

of available data and encompasses the principal physical non-financial commodities that are the subject of active, liquid futures markets.* The liquidity constraint is aimed at promoting cost-effective implementation and true investability.

The S&P GSCI® index is largely diversified across a broad range of commodity sectors, which enables a high level of diversification, both across sub-sectors and within each sub-sector. The diversification also reduces the impact of idiosyncratic events, which significantly impact individual commodity markets, but are muffled when aggregated to the level of S&P GSCI®. The index is seen as a good indicator of general price movements and inflation in the world economy by market participants and researchers. The S&P GSCI total return reflects the performance of total return investment in commodities and is computed as the product of three variables: (i) the value of S&P GSCI on the preceding S&P GSCI Business Day, (ii) one plus the sum of the Contract Daily Return and the Treasury Bill Return on the S&P GSCI Business Day on which the calculation is made, and (iii) one plus the Treasury Bill Return for each non-S&P GSCI Business Day since the preceding S&P GSCI Business Day.†

As indicated above, the computation of the index recognizes that idiosyncratic events may significantly impact on individual commodity sectors. Jensen *et al.* (2000) also show that commodity prices are sensitive to economic conditions and are likely to vary across commodities. Consequently, we do not rely exclusively on the S&P GSCI® composite index and investigate the performance of 10 sub-indexes as well: S&P GSCI® Energy, S&P GSCI® Light-Energy, S&P GSCI® Non-Energy, S&P GSCI® Reduced Energy, S&P GSCI® Agriculture, S&P GSCI® Livestock, S&P GSCI® Petroleum, S&P GSCI® Industrial Metal, S&P GSCI® Precious Metals and S&P GSCI® Softs. The sub-indexes of the S&P GSCI® are constructed using the same rules regarding world production weights, the methodology for rolling and other functional characteristics as disclosed in the S&P GSCI® manual.

The stock market returns data used in this paper are drawn from the S&P 500 index which includes the 500 leading companies in leading industries of the U.S. economy. The index only includes large companies with market capitalization in excess of US\$3.5 billion and covers approximately 75% of U.S. equities. Given that it includes a significant segment of the total value of the equity market, investors, academics and market participants alike consider the index to be an ideal proxy for the total market return. The components of the S&P 500 index are classified according to the Global Industry Classification Standard (GICS). All data used in this study is sourced from Datastream.

18.4.2 Data Description

Table 18.1 presents descriptive statistics for the S&P 500 index and each S&P GSCI® index used in this paper. Between January 1999 and December 2009, the S&P 500 posted a maximum (minimum) return of 0.102 (−0.164). The commodities total return index, S&P GSCI®, shows a maximum (minimum) return of 0.123 (−0.131) for the same period. The mean return of 0.00141 produced by the equity index for the sample period is superior to that obtained by the S&P GSCI®, which achieved a mean return of 0.00084.

* See the following link for further details on the components, weights, and construction of the S&P GSCI® index: <http://www.standardandpoors.com/servlet/BlobServer?blobheadername3=MDT-Type&blobcol=urldata&blobtable=MungoBlobs&blobheadername2=inline%3B+filename%3D&blobkey=id&blobheadername1=content-type&blobwhere=1243748755917&blobheadername3=UTF-8>.

† To reflect the performance of a total return investment in commodities, four separate but related indexes have been developed based on the S&P GSCI: (1) the S&P GSCI Spot Index, which is based on the price levels of the contracts included in the S&P GSCI; (2) the S&P GSCI Excess Return Index (S&P GSCI ER), which incorporates the returns of the S&P GSCI Spot Index as well as the discount or premium obtained by ‘rolling’ hypothetical positions in such contracts forward as they approach delivery; (3) the S&P GSCI Total Return Index (S&P GSCI TR), which incorporates the returns of the S&P GSCI ER and interest earned on hypothetical fully collateralized contract positions on the commodities included in the S&P GSCI; and (4) the S&P GSCI Futures Price Index (S&P GSFPI), which is intended to serve as a benchmark for the fair value of the futures contracts on the S&P GSCI traded on the Chicago Mercantile Exchange (CME).

TABLE 18.1 Descriptive Statistics for the S&P 500 Stock Index and the S&P GSCI Commodity Index Returns From January 1999 Through December 2009 for a Total of 767 Weekly Observations

Market	Mean	Median	Std Dev	Skewness	Kurtosis	Min	Max
S&P 500	0.00141	0.00326	0.0246	4.327	-0.582	-0.164	0.102
Commodity	0.00084	0.00137	0.0313	1.698	-0.433	-0.131	0.123
Energy	0.00145	0.00364	0.0439	1.076	-0.339	-0.186	0.173
Petroleum	0.00213	0.00447	0.0450	1.751	-0.402	-0.213	0.207
Non-energy	-0.00004	0.00064	0.0183	3.414	-0.497	-0.094	0.073
Reduced energy	0.00059	0.00127	0.0263	2.344	-0.541	-0.119	0.102
Light energy	0.00036	0.00135	0.0221	2.969	-0.614	-0.106	0.092
Industrial metal	0.00089	0.00094	0.0287	3.662	-0.047	-0.133	0.147
Precious metals	0.00141	0.00077	0.0239	3.926	-0.031	-0.127	0.119
Agriculture	-0.00072	-0.00006	0.0264	2.056	-0.127	-0.131	0.115
Livestock	-0.00070	-0.00053	0.0195	2.241	-0.610	-0.119	0.057
Softs	-0.00029	0.00012	0.0277	1.370	-0.060	-0.124	0.112

As Table 18.1 shows, two S&P GSCI® sub-indexes, energy and petroleum, outperformed the S&P 500 index. The sub-indexes for non-energy, agriculture, livestock and softs posted negative mean returns for the same period. The return distribution for all market indexes seems to be non-normal and shows positive skewness. The S&P GSCI® appears to be more volatile than the S&P 500.

Table 18.2 presents the pairwise correlation of returns for all pairs of market indexes used in the study. The S&P 500 is positively correlated to the S&P GSCI®, with a correlation coefficient of 0.165, as well as to all the sub-indexes for the commodity index. All the correlations between the S&P 500 and the S&P GSCI commodity indexes (the main and sub-indexes) are below 0.3, indicating low comovements of asset returns. Table 18.2 also sheds light on the variable correlation of returns between that of the main commodity index, S&P GSCI® and its sub-indexes. Some sub-indexes, e.g. energy, petroleum, reduced energy and light energy, exhibit high correlations (above 0.9) with the main index. Others have correlation coefficients below 0.5 (industrial metals, precious metals, agriculture, livestock, softs) with the main index. S&P GSCI® livestock is shown to have the lowest correlation with the main index. From Table 18.2 we can also note that S&P GSCI® livestock exhibits the lowest correlations for all pairs of sub-indexes.

18.5 Empirical Analysis of Comovements

In this section we present evidence on the wavelet squared coherency for the S&P 500 and the S&P GSCI® commodity index as well as that between S&P 500 and 10 sub-indexes of the S&P GSCI® commodity index. As indicated above, the wavelet squared coherency analysis involves two time-series given as input in a time–frequency space. Consequently, the wavelet squared coherency is presented using contour plots as it involves three dimensions, frequency, time and wavelet squared coherency power (similar to Rua and Nunes (2009)).

Frequency and time are represented on the vertical and horizontal axes, respectively. Frequency is converted into time units (years) to simplify interpretation and ranges from the highest frequency of one week (top of the plot) to the lowest frequency of four years (bottom of the plot). The wavelet squared coherency is depicted by a grey scale in the figures. The scale for the wavelet coherency is interpreted in terms of the darkness of the grey colour, where an increasing value of the wavelet squared coherency corresponds to a deepening darkness of grey and imitates the height in a surface plot. Subsequently, through a visual assessment of the graphs, we can identify both the time intervals (on the horizontal axis) and frequency bands (on the vertical axis) where the two markets (commodity and stock) move together.

The frequency scale allows us to differentiate the effects of comovements between the short term and the long term. We consider short term to relate to frequencies ranging from two weeks to six months.

TABLE 18.2 Pearson's Correlation Coefficients for the S&P 500 Index and the S&P GSCI Commodity Index Returns. *, **, *** Denote Statistical Significance at the 5%, 1% and 0.1% Level, Respectively

	S&P 500	Commodity	Energy	Petroleum	Non-energy	Reduced energy	Light energy	Industrial metals	Precious metals	Agriculture	Livestock
Commodity	0.165***										
Energy	0.133***	0.973***									
Petroleum	0.137***	0.933***	0.955***								
Non-energy	0.232***	0.529***	0.351***	0.345***							
Reduced energy	0.184***	0.981***	0.922***	0.885***	0.645***						
Light energy	0.209***	0.927***	0.828***	0.797***	0.788***	0.978***					
Industrial metals	0.264***	0.410***	0.304***	0.308***	0.632***	0.470***	0.547***				
Precious metals	0.057	0.309***	0.237***	0.244***	0.448***	0.349***	0.400***	0.348***			
Agriculture	0.136***	0.409***	0.255***	0.245***	0.861***	0.517***	0.648***	0.258***	0.256***		
Livestock	0.114***	0.193***	0.124***	0.119***	0.332***	0.233***	0.278***	0.094***	0.046	0.090**	
Softs	0.141***	0.328***	0.232***	0.238***	0.549***	0.389***	0.462***	0.261***	0.241***	0.588***	0.015

Long term can be thought of as investing in frequencies ranging from two to three years. In this unified framework, a dark grey area at the bottom (top) of the figures would indicate strong comovements at low (high) frequencies, whereas a dark grey area on the left-hand (right-hand) side signifies strong comovements at the start (end) of the sample period.

The thick black continuous line in the figures sequesters regions where the wavelet squared coherency is statistically significant at the 5% level, i.e. the wavelet squared coherency is statistically significant within such a delimited time–frequency area. The significance level was simulated as explained in [Section 18.3](#) using a Monte Carlo method of 10,000 sets of two white noise time-series with the same length as the series under analysis.

Figure 18.1 depicts the standard wavelet squared coherency between the S&P 500 and the S&P GSCI® commodity index and reveals several interesting features. First, we find no evidence of comovements of the S&P 500 total return and the S&P GSCI® commodity index total return in the higher frequencies, indicating diversification benefits in the short term. The onset of the financial crisis post-2007 does not alter this short-term diversification benefit. Secondly, the comovement of returns between the two indexes at low frequencies before 2007 is very weak. The conclusion that may be drawn here is that long-term investors derived diversification benefits by combining the two assets in a portfolio during this time period. There is, however, some notable evidence of increasing comovements at low, rising to medium, frequencies towards the end of the series (after 2007). This perceptible rise in market dependencies between the two return index series coincides with the onset of the most recent financial crisis. As has been noted in the literature, in periods of deleveraging, correlations between uncorrelated assets




FIGURE 18.1 Comovement of the S&P 500 stock index and the S&P GSCI commodity index. This figure presents the wavelet squared coherency between the S&P 500 index returns and the S&P GSCI commodity index returns. Time and frequency are represented on the horizontal and vertical axes, respectively. Frequency is converted into years. The wavelet squared coherency is represented by the grey scale where increasing darkness of the grey scale corresponds to an increasing value and mimics the height in a surface plot. The heavy black line isolates the statistically significant area at the 5% significance level.

can increase to +1. This appears to be the case with the long-term fluctuations of the two financial assets return series. Thus the evidence implies that including commodities in an equity portfolio may not provide significant diversification benefits for long-term investors during a period of financial crisis.

Given the possible variation in sensitivity of commodity prices to economic conditions, we further extend our empirical investigations by examining the comovement between the S&P 500 total return and the sub-indexes of the S&P GSCI® commodity total returns. Figure 18.2 reports the results of the examination of the comovements between the S&P 500 total return and three sub-indexes of the S&P GSCI® commodity index total return: S&P GSCI® Energy Total Return, S&P GSCI® Light Energy Total Return and S&P GSCI® Red Energy Total Return. The results mirror the finding for the comovements between the S&P 500 total return and the S&P GSCI® commodity index total return discussed above. Once again, a clear distinction can be made between the short-term and long-term diversification benefits of combining the two assets. We also find weak dependencies between the S&P 500 total return and the S&P GSCI® Energy total return, the S&P 500 total return and the S&P 500 GSCI® Red Energy total return and the S&P 500 total return and the S&P 500 GSCI® Light Energy total return at higher frequencies for the whole data period. Analogous to the results discussed above, there is also evidence of rising comovement between the S&P 500 and the S&P GSCI® sub-indexes after 2007. Thus, the short-term diversification benefits across time of combining equity and commodities extend to these sub-indexes of S&P 500 GSCI®.

Figure 18.3 shows the result of examining the comovements between the S&P 500 total return and the S&P GSCI® non-Energy total return. We find no consistent evidence of comovement between the two indexes in the short term. Parallel to the documented relationship between the S&P 500 and S&P GSCI® commodity index above, we find evidence of increasing levels of long-term comovements from low, rising to intermediate, frequencies towards the end of the series.

The comovement of returns between the S&P 500 total return and the S&P GSCI® industrial metals total return and that of the S&P 500 and S&P GSCI® precious metals total return reveals some important details (Figure 18.4). The short-term and long-term evidence regarding the comovement between the S&P GSCI® precious metals total return is very weak at best. Thus this segment of the S&P GSCI® index provides significant diversification benefits for both long-term and short-term investors. The evidence only suggests strong short-term diversification benefits when S&P 500 and S&P GSCI® industrial metals are considered.

Examining the comovement of returns between the S&P 500 total return and the S&P GSCI® agriculture sub-indexes, we find weak comovement between the S&P 500 total and the S&P GSCI® Agriculture total return, both in the short and long term (see Figure 18.5). The evidence, however, indicates relatively stronger comovement between the two indexes towards the end of the series. This result is replicated for




FIGURE 18.2 Comovement of the S&P 500 index and the S&P GSCI Energy sub-indexes. This figure presents the wavelet squared coherency between the S&P 500 index returns and three S&P GSCI Energy indexes: the S&P GSCI Energy Total returns, the S&P GSCI® Red Energy Total Return and the S&P GSCI® Light Energy Total Return. Time and frequency are represented on the horizontal and vertical axes, respectively. Frequency is converted into years. The wavelet squared coherency is represented by the grey scale where increasing darkness of the grey scale corresponds to an increasing value and mimics the height in a surface plot. The heavy black line isolates the statistically significant area at the 5% significance level.




FIGURE 18.3 Comovement of the S&P 500 index and the S&P GSCI non-Energy sub-index. This figure presents the wavelet squared coherency between the S&P 500 index returns and the S&P GSCI non-Energy index. Time and frequency are represented on the horizontal and vertical axes, respectively. Frequency is converted into years. The wavelet squared coherency is represented by the grey scale where increasing darkness of the grey scale corresponds to an increasing value and mimics the height in a surface plot. The heavy black line isolates the statistically significant area at the 5% significance level.




FIGURE 18.4 Comovement of the S&P 500 index and the S&P GSCI Metals sub-indexes. This figure presents the wavelet squared coherency between the S&P 500 index returns and the S&P GSCI Industrial Metals index and the S&P 500 index returns and the S&P GSCI Precious Metals index. Time and frequency are represented on the horizontal and vertical axes, respectively. Frequency is converted into years. The wavelet squared coherency is represented by the grey scale where increasing darkness of the grey scale corresponds to an increasing value and mimics the height in a surface plot. The heavy black line isolates the statistically significant area at the 5% significance level.




FIGURE 18.5 Comovement of the S&P 500 index and three S&P GSCI sub-indexes. This figure presents the wavelet squared coherency between the S&P 500 index returns and three S&P GSCI sub-indexes: the S&P GSCI Agricultural Total Return, the S&P GSCI Livestock Total Return and the S&P GSCI Softs Total Return. Time and frequency are represented on the horizontal and vertical axes, respectively. Frequency is converted into years. The wavelet squared coherency is represented by the grey scale where increasing darkness of the grey scale corresponds to an increasing value and mimics the height in a surface plot. The heavy black line isolates the statistically significant area at the 5% significance level.

the comovement of returns between the S&P 500 total return and the S&P GSCI® Livestock total return. We also find that combining an equity portfolio and S&P GSCI® Softs provides important diversification benefits for both the short-term and long-term investors (see Figure 18.5). The onset of the financial crisis does not seem to impact on these benefits.

18.6 Conclusion

The diversification properties of commodities have long been understood by market participants. The recent financial crisis has, however, raised apprehension that the diversification line of reasoning in favour of commodities investment is not robust. This brings to the forefront the question of the independence of commodity index returns. This study therefore investigates the comovements of returns of the S&P 500 and the S&P GSCI Commodity Index using an enhanced methodology, the three-dimensional analysis of wavelet coherency, during the period January 1999 to December 2009. The sample periods enable us to compare the pre-financial crisis era comovements with that of the onset of the crisis. The wavelet methodology represents a refinement and allows the examination of the time-and-frequency varying comovements of stock markets within a unified framework and enables a distinction between short-term and long-term benefits of diversification.

Generally, we find weak evidence of comovement between the S&P 500 total return and the S&P GSCI® Commodity Index total return for our sample period. This implies that commodities still have diversification value when combined with equities during the recent financial crisis. The argument for diversification benefits is even stronger when we segment our analyses and examine the gains from the point of view of short-term and long-term investors. From the point view of short-term investors, the onset of the current financial crisis has not dampened the diversification value of commodities. The pre-financial crisis period and the onset of the financial crisis at the end of 2007 show a similar lack of correlation at high frequencies. From the point of view of the long-term investor, however, we note a rising comovement of returns for the two indexes after 2007, which coincides with the onset of the financial crisis. It should be noted, though, that the comovement is not strong. Thus, a case for limited long-term diversification gains could be made during the financial crisis.

Earlier studies alluded to the possible variation in the sensitivity of commodity prices to economic conditions across commodities. As a consequence, we further investigate the comovement of the S&P 500 and 10 sub-indexes of the S&P GSCI index. This empirical exercise generally confirms our earlier findings. For the short-term investor, the current financial crisis shows no measurable impacts on the diversification benefits of combining equity and commodities. From the point of view of the long-term

investor, typified by fluctuations at low frequencies, the evidence suggests an increase in comovement at the onset of the financial crisis. The evidence, however, suggests weak correlations. Thus there may be limited benefits even in time of crisis. There are some important distinctions that can be made for combining equity and the various segmented commodity indexes. For instance, the diversification benefits of combining equity (S&P 500 total return) and the S&P GSCI® Precious Metal total return is robust for both long-term and short-term investors across time.

Future research may construct different portfolios (with and without commodities) based on wavelet analysis (high-versus low-dependence portfolios) and compare their realized risk-adjusted performance.

Acknowledgements

We would like to thank the anonymous reviewers and the editor of the journal for their insightful comments. The financial support of the Academy of Finland (project 132913) is also gratefully acknowledged.

References

- Abanomey, W.S. and Mathur, I., The hedging benefits of commodity futures in international portfolio construction. *J. Altern. Invest.*, 1999, **2**, 51–62.
- Adel, B. and Salma, J., The Greek financial crisis, extreme comovement, and contagion effects in the EMU: a copula approach. *Int. J. Account. Financ. Rep.*, 2012, **2**, 289–307.
- Akey, R.P., Commodity: a case for active investment. *J. Altern. Invest.*, 2005, **8**, 8–29.
- Aloui, R., Ben Aïssa, M.S. and Nguyen, D.K., Global financial crisis, extreme interdependences, and contagion effects: the role of economic structure? *J. Bank. Finance*, 2011, **35**, 130–141.
- Ang, A. and Chen, J., Asymmetric correlations of equity portfolios. *J. Financ. Econ.*, 2002, **63**, 443–494.
- Ankrum, E.M. and Hensel, C.R., Commodities in asset allocation: a real asset alternative to real estate? *Financ. Anal. J.*, 1993, **49**, 20–29.
- Anson, M., Maximizing utility with commodity futures diversification. *J. Portfol. Mgmt*, 1999, **25**, 86–94.
- Bartram, S.M. and Bodnar, G.M., No place to hide: the global crisis in equity markets in 2008/2009. *J. Int. Money Finance*, 2009, **28**, 1246–1292.
- Belousova, J. and Dorfleitner, G., On the diversification benefits of commodities from the perspective of euro investors. *J. Banking & Finan.*, 2012, **36**, 2455–2472.
- Billingsley, R.S. and Chance, D.M., Benefits and limitations of diversification among commodity trading advisors. *J. Portfol. Mgmt*, 1996, **23**, 65–80.
- Candelon, B., Piplack, J. and Straetmans, S., On measuring synchronization of bulls and bears: the case of East Asia. *J. Bank. Finance*, 2008, **32**, 1022–1035.
- Chance, D.M., Managed futures and their role in investment portfolios. The Research Foundation of the Institute of Chartered Financial Analysts. Charlottesville, VA, 1994.
- Cheung, C.S. and Miu, P., Diversification benefits of commodity futures. *J. Int. Financ. Mkts, Inst. & Money*, 2010, **20**, 451–474.
- Claessens, S., Kose, M.A. and Terrones, M.E., The global financial crisis: how similar? How different? How costly? *J. Asian Econ.*, 2010, **21**, 247–264.
- Conover, C.M., Jensen, G.R., Johnson, R.R. and Mercer, J.M., Is now the time to add commodities to your portfolio? *J. Invest.*, 2010, **19**, 10–19.
- Daskalaki, C. and Skiadopoulos, G.S., Should investors include commodities in their portfolios after all? New evidence. *J. Banking & Finan.*, 2011, **35**, 2606–2626.
- Daubechies, I., Ten lectures on wavelets (CBMS-NSF Regional Conference Series in Applied Mathematics). Society for Industrial & Applied Mathematics, 1992.
- Dwyer, G.P. and Tkac, P., The financial crisis of 2008 in fixed income markets. *J. Int. Money Finance*, 2009, **28**, 1293–1316.
- Edwards, F.R. and Liew, J., Managed commodity funds. *J. Futures Markets*, 1999, **19**, 377–411.

- Erb, C.B. and Harvey, C.R., The tactical and strategic value of commodity futures. *Financ. Anal. J.*, 2006, **62**, 69–97.
- Fortenbery, T.R. and Hauser, R.J., Investment potential of agricultural futures contracts. *Am. J. Agric. Econ.*, 1990, **72**, 721–726.
- Garcia, R. and Tsafack, G., Dependence structure and extreme comovement in international equity and bond markets. *J. Bank. Finance*, 2009, **8**, 1954–1970.
- Garrett, I. and Taylor, N., Intraday and interday basis dynamics: evidence from the FTSE100 Index Futures Market. *Stud. Nonlin. Dynam. Econometr.*, 2001, **5**, 133–152.
- Geman, H., *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*, 2005 (Wiley: Chichester).
- Gorton, G., Hayashi, F. and Rouwenhorst, K.G., Commodity futures: a Japanese perspective. Yale ICF Working Paper No. 05–27, 2005.
- Gorton, G. and Rouwenhorst, K.G., Facts and fantasies about commodity futures. *Financ. Anal. J.*, 2006, **62**, 47–68.
- Gorton, G. and Rouwenhorst, K.G., Are commodity futures too risky for your portfolio? Hogwash! Knowledge Wharton, 2006b.
- Graham, M. and Nikkinen, J., Co-movement of the Finnish and international stock markets: a wavelet analysis. *Eur. J. Finance*, 2011, **17**, 409–425.
- Graham, M., Kiviaho, J. and Nikkinen, J., Integration of 22 emerging stock markets: a three dimensional analysis. *Glob. Finance J.*, 2012, **23**, 34–47.
- Greer, R.J., Methods for institutional investment in commodity futures. *J. Deriv.*, 1994, **2**, 28–36.
- Grinsted, A., Moore, J.C. and Jevrejeva, S., Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlin. Process. Geophys.*, 2004, **11**(5/6), 561–566.
- Ibbotson Associates, *Strategic Asset Allocation and Commodities*, 2006 (Ibbotson Associates: Chicago, IL).
- Idzorek, T.M., Commodities and strategic asset allocation. In *Intelligent Commodity Investing*, edited by H. Till and J. Eagleeye, 2007 (Risk Books: London).
- Jensen, G., Johnson, R.R. and Mercer, J.M., Efficient use of commodity futures in diversified portfolios. *J. Fut. Mkts*, 2000, **5**, 489–506.
- Jensen, G., Johnson, R.R. and Mercer, J.M., Tactical asset allocation and commodity futures. *J. Portfol. Mgmt*, 2002, Summer, 100–110.
- Kaplan, P.D. and Lummer, S.L., Update: GCSI collateralized futures as a hedging and diversification tool for institutional portfolios. *J. Invest.*, 1998, Winter, 11–17.
- Kenourgios, D. and Padhi, P., Emerging markets and financial crises: regional, global or isolated shocks? *J. Multinatl. Financ. Mgmt*, 2012, **22**, 24–38.
- Kyle, A.S. and Xiong, W., Contagion as a wealth effect. *J. Finance*, 2001, **56**, 1401–1440.
- Longin, F. and Solnik, B., Extreme correlation of international equity markets. *J. Finance*, 2001, **56**, 649–676.
- Lummer, S.L. and Siegel, L.B., GSCI collateralized futures: a hedging and diversification tool for institutional portfolios. *J. Invest.*, 1993, **2**, 75–82.
- Mallat, S., *A Wavelet Tour of Signal Processing*, 1998 (Academic Press: San Diego, CA).
- Markowitz, H.M., Portfolio selection. *J. Finance*, 1952, **7**, 77–91.
- Merath, J., Commodities are part of an optimized portfolio. Global Trends, Credit Suisse, 2006.
- Nguyen, V.T.T. and Sercu, P.M.F.A., Tactical asset allocation with commodity futures: implications of business cycle and monetary policy (May 26, 2010). In *Paris December 2010 Finance Meeting EUROFIDAI – AFFI*, 2010. Available online at SSRN: <http://ssrn.com/abstract=1695889>.
- Nikkinen, J., Pynnönen, S., Ranta, M. and Vähämaa, S., Crossdynamics of exchange rate expectations: a wavelet analysis. *Int. J. Finance Econ.*, 2011, **16**, 205–217.
- Patten, A., On the out of sample importance of skewness and asymmetric dependence for asset allocation. *J. Financ. Econometr.*, 2004, **2**, 130–168.

- Rua, A. and Nunes, L.C., International comovement of stock returns: a wavelet analysis. *J. Empir. Finance*, 2009, **11**, 632–639.
- Satyanarayan, S. and Varangis, P., Diversification benefits of commodity assets in global portfolios. *J. Invest.*, 1996, **5**, 69–78.
- Silvennoinen, A. and Thorp, S., Financialization, crisis and commodity correlation dynamics. Working Paper, Queensland University of Technology, 2010.
- Tang, K. and Xiong, W., Index investing and the financialization of commodities. Working Paper, Princeton University, 2010.
- Torrence, C. and Compo, G.P., A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.*, 1998, **79**, 61–78.
- Yamori, N., Co-movement between commodity market and equity market: does commodity market change? *Mod. Econ.*, 2011, **2**, 335–339.

Portfolio Selection with Commodities Under Conditional Copulas and Skew Preferences

Carlos González-
Pedraz

Manuel Moreno

Juan Ignacio Peña

19.1	Introduction	404
19.2	Portfolio Choice with Commodity Futures and Skewness.....	405
19.3	Multivariate Conditional Copula with Asymmetry	407
	Modelling Univariate Processes • Modelling Copula Functions • Model Estimation and Portfolio Optimization	
19.4	Empirical Application.....	412
	Data and Preliminary Analysis • Estimation of the Conditional Copula Model • Optimal Portfolio Results	
19.5	Conclusion	427
	Acknowledgements	428
	Supplemental Data	428
	References.....	428
	Appendix A: Copula Functions.....	430
	Appendix B: The Two-Stage Log-Likelihood Function.....	432

This article investigates the portfolio selection problem of an investor with three-moment preferences taking positions in commodity futures. To model the asset returns, we propose a conditional asymmetric t copula with skewed and fat-tailed marginal distributions, such that we can capture the impact on optimal portfolios of time-varying moments, state-dependent correlations and tail and asymmetric dependence. In the empirical application with oil, gold and equity data from 1990 to 2010, the conditional t copula portfolios achieve better performance than those based on more conventional strategies. The specification of higher moments in the marginal distributions and the type of tail dependence in the copula has significant implications for out-of-sample portfolio performance.

Keywords: Portfolio selection; Commodity futures; Conditional copulas; Skew preferences

JEL Classification: C46, G11, G13

19.1 Introduction

Financial investors mainly take positions in commodity futures contracts as a natural way to gain exposure to commodity risk without owning the physical asset. Erb and Harvey (2006) and Gorton and Rouwenhorst (2006) find that historically, commodity futures exhibited little comovement, zero or even negative correlation with stock returns, with Sharpe ratios fairly close to those of equities. Therefore, according to traditional portfolio theory, commodities should increase diversification when included in equity portfolios and may help enhance the portfolio's risk-return profile. Possibly boosted by the potential for such diversification benefits, investments in commodity futures indexes and related instruments grew quickly after the early 2000s* (see Büyüksahin and Robe (2014), Etula (2010), Hong and Yogo (2012) and Tang and Xiong (2012) for some analysis of this recent boom).

Despite the growing interest in commodities as investment vehicles, few studies have analysed the optimal portfolio allocation taking into account the stylized features of commodity futures. A standard mean-variance framework might not be appropriate for portfolios that contain commodity futures due to their returns' specific distributional characteristics, such as the presence of serial correlation, heavy tails and skewness (Gorton and Rouwenhorst 2006, Kat and Oomen 2007, Boerger *et al.* 2009, Ohana 2010, Daskalaki and Skiadopoulos 2011, Zolotko and Okhrin 2012). Instead, we propose a more flexible model to be used in the portfolio selection problem of an equity investor when cash-collateralized commodity futures are part of the investment opportunity set. Our approach combines a three-moment preference specification with time-varying multivariate density models that describe the statistical properties of commodities and equity returns, as well as their interactions.

With respect to the investor's preferences, we consider an allocation problem in which the investor's objective function is determined by the mean, variance and skewness of portfolio returns (similar to Guidolin and Timmermann (2008), Harvey *et al.* (2010) and Jondeau and Rockinger (2012)). With fairly general assumptions, investors show a preference for positive skewness in return distributions and aversion to downside risk (negative skewness). That is, in our proposed three-moment preference specification, the investor is eager to decrease the chance of large negative deviations, which could reduce the final value of the portfolio. Expanding the standard mean-variance set-up by including the third moment of the portfolio returns has been investigated for traditional assets, such as stocks.[†] As far as we know, however, this approach has not yet been analysed for a portfolio choice problem with commodities, for which skewness seems likely to play a role due to the specific features of commodity assets. For instance, the possibility of shortages in supply may produce jumps in prices, leading to skewness in the returns of futures contracts.

Regarding the multivariate density model, we offer a flexible approach to specifying the joint distribution of returns using conditional copula models. Copula functions help disentangle the particular characteristics of the univariate distributions of equity and commodity returns from their dependence structure. We combine conditional copula theory, as presented in Patton (2006a, b), with the implicit copula functions of multivariate normal mixtures defined by Demarta and McNeil (2005) and Embrechts *et al.* (2003). As our most general model, we propose a conditional skewed *t* copula with generalized Student's *t* marginal distributions. This copula model allows for asymmetric and tail dependence in a multivariate framework, and includes symmetric and linear dependence as special cases. Furthermore, the conditional set-up enables us to capture time-varying investment opportunities through time-varying moments and changes in the dependence parameters. These copula models are

* Investments in commodity index funds increased from around \$50 billion, at the end of 2004 to a peak of \$200 billion in 2008; after a drop during the recession, they increased again to a second peak of around \$300 billion at the end of the third quarter of 2010. See Irwin and Sanders (2011).

† Some early works on how skewness affects portfolio selection include Samuelson (1970) and Kraus and Litzenberger (1976). Harvey and Siddique (2000) build on these ideas to provide an empirical test of the effect of co-skewness on asset prices. Barberis and Huang (2008) and Mitton and Vorkink (2007) also suggest, from different perspectives, that the skewness of individual assets may also influence investors' portfolio decisions.

particularly easy to sample from, and therefore, we opt for solving the optimization problem numerically using Monte Carlo simulations (as in Patton (2004) and Ortobelli *et al.* (2010)).

We apply our theoretical approach to weekly data of crude oil and gold futures and the S&P 500 equity index, for the period from June 1990 to September 2010, reserving the observations from September 2006 to September 2010 for an out-of-sample performance evaluation. We examine four primary issues:

1. Is there asymmetric and tail dependence among commodity and equity returns?
2. Are there discrepancies in the optimal portfolio allocations between our conditional copula approach and other more traditional benchmarks, such as the equally weighted or Gaussian strategies?
3. Do these discrepancies translate into economically relevant performance differences among methods?
4. Is there a single key factor explaining these discrepancies?

First, our preliminary statistics and in-sample and out-of-sample results show evidence in favour of heavy tails and skewness in the univariate behaviour and extreme and asymmetric dependence amongst oil, gold and equity. Second, we also uncover substantial discrepancies between portfolio optimal weights of conditional t copulas and the portfolio weights provided by more conventional alternatives, especially for more aggressive investors and when there are no restrictions on short-selling positions in equity. Third, in most cases, the discrepancies in portfolio weights translate into economically more profitable investment ratios and better relative performance measures with respect to the alternative procedures. Depending on the investor's preference specification, the gains of considering the conditional copula model with tail and asymmetric dependence instead of the equally weighted portfolio represent up to 86 basis points per year for the period 2006–2010. When variance and loss aversion increase, portfolio strategies based on more flexible copulas are less likely to produce large performance differences. Fourth, no single factor offers a sufficient explanation of these differences. Rather, we find various explanatory elements, including, the specification of the univariate processes, in terms of conditional volatility, skewness and fat tails; and the presence of tail and asymmetric dependence.

The remainder of this article is organized as follows. Section 19.2 formulates the investor's objective function and the portfolio choice problem. In Section 19.3, we describe the multivariate conditional copula model, the estimation methodology and the numerical implementation. Section 19.4 presents the empirical application: the data description and preliminary statistics, the in-sample estimation and the out-of-sample portfolio allocation results. We conclude in Section 19.5. In Appendices A and B, we develop the explicit functional forms of the copula densities considered and the joint log-likelihood function. Furthermore, in a separate supplemental data section, we provide more technical details and additional tables and figures.

19.2 Portfolio Choice with Commodity Futures and Skewness

In this section, we present the portfolio selection problem of an investor with mean-variance-skewness preferences that takes positions in commodity futures and other risky spot contracts, such as stocks.

No money changes hands when futures are bought or sold; just a margin is posted to running settle gains and losses. Without any upfront payment, it is not clear how to define the rate of return (Dusak 1973). Following the common approach in the literature to analyse commodity futures as an asset class (Gorton and Rouwenhorst 2006, Hong and Yogo 2012), we assume that long and short positions are fully collateralized. That is, the initial margin deposit corresponds with the overall notional value of the futures contract and indicates the initial capital investment related to that position (long or short). In this way, we can control for the leverage involved in futures positions, and we can make fair comparisons with spot contracts.* Therefore, taking collateral in futures contracts into account would affect the computation of their rates of return and the budget constraint of the investor's problem, as we will see.

* This assumption can also be relaxed, and smaller fractions of the nominal value can be considered in the problem set-up.

Formally, our portfolio choice problem can be formulated in terms of an investor who maximizes expected utility at period $t + 1$ by building at time t a portfolio that includes two groups of assets: a group with n commodity futures contracts, and another group with $N - n$ spot contracts, such as stocks. Taking the collateral into account, the gross return of a position in the commodity futures contract i at time $t + 1$ is given by

$$(1 + R_{i,t+1}) = \frac{S_{i,t+1}}{S_{i,t}} (1 + R_{t+1}^f), \quad i = N - n + 1, \dots, N, \quad (19.1)$$

where $S_{i,t}$ and $S_{i,t+h}$ are the futures settlement prices at times t and $t + 1$, respectively, and $(1 + R_{t+1}^f)$ is the gross return on cash over the period, or the interest earned on the initial margin deposit (Hong and Yogo 2012). For this set of N investment opportunities, the wealth at time $t + 1$ equals the gross return of the portfolio over the period, $1 + R_{t+1}(\omega_t)$, defined as

$$1 + R_{t+1}(\omega_t) = 1 + \sum_{j=1}^N \omega_t^j (\exp(r_{j,t+1}) - 1), \quad (19.2)$$

where $\omega_t = (\omega_t^1, \dots, \omega_t^{N-n}, \omega_t^{N-n+1}, \dots, \omega_t^N)'$ is the vector of portfolio weights (for spot and futures contracts), chosen at time t , and $r_{j,t+1} = \log(1 + R_{j,t+1})$ is the continuously compounded return of asset j over the period.

As in Etula (2010), using this distinction between commodity futures contracts and traditional securities, we are considering that investors trade only marketable securities and avoid trading the physical underlying. That is, the investor has a traditional equity portfolio and gains exposure to commodity risk taking positions in futures contracts. In addition, as in Hong and Yogo (2012) and Gorton and Rouwenhorst (2006), futures are assumed to be non-zero-investment positions and their collateral is taken into account.

As is well known, returns on financial assets generally deviate from the Gaussian distribution, displaying heavy tails and skewness. This departure from normality is even greater in the case of commodities, magnified by the well-documented presence of positive and negative spikes in the data-generating process of commodity returns (see e.g. Cartea and Figueroa (2005) and Casassus and Collin-Dufresne (2005), among others). The fundamentals underlying commodity price formation are key determinants of these statistical properties. Accordingly, the presence of jumps can be explained by the convex relationship between commodity prices and the balance among supply, inventories and demand (see Routledge *et al.* (2000)). Thus, adding commodity assets to traditional portfolios will constitute a significant source of skewness for the portfolio's returns, increasing the importance of considering the third moment in the portfolio selection problem.

For that reason, in our approach the investor's objective consists of choosing a wealth allocation $\omega_t = (\omega_t^1, \dots, \omega_t^N)$ that maximizes the expected portfolio return penalized for the variance and negative skewness of the portfolio returns. That is, for each time t , the optimal weights are given by

$$\begin{aligned} \omega_t^* = \arg \max_{\omega_t \in \mathcal{D}} & \left(\mathbb{E}_t[R_{t+1}(\omega_t)] - \varphi_V \text{Var}_t[R_{t+1}(\omega_t)] \right. \\ & \left. + \varphi_S \text{Skew}_t[R_{t+1}(\omega_t)] \right), \end{aligned} \quad (19.3)$$

where $\mathbb{E}_t(\cdot)$, $\text{Var}_t(\cdot)$ and $\text{Skew}_t(\cdot)$ are the first three moments of the portfolio returns conditioned on the information set $\mathcal{F}_t$ available at time t . The parameters $\varphi_V \geq 0$ and $\varphi_S \geq 0$ determine the impact of variance (traditional risk aversion) and skewness (loss aversion) on the investor's utility. By adding aversion to negative skewness, we acknowledge the possibility that an investor might accept a lower expected return if there is a chance of high positive skewness, such as in the form of a large probability of positive jumps.

The importance of skewness in asset pricing and portfolio choice has been studied for traditional assets by Kraus and Litzenberger (1976), Harvey and Siddique (2000) and Harvey *et al.* (2010), among others. Apart from analysing skewness, other studies also consider the impact of kurtosis of the portfolio

returns, see for example Guidolin and Timmermann (2008). Using our specification of the investor's objective function, we want to focus specifically on the effect of skewness and asymmetric dependence in the optimal portfolio choice when commodities, which are characterized by displaying significant skewness, are part of the investment opportunity set. For that reason, we do not include additional higher moments in our specification. In spite of that, we do not deny that the inclusion of kurtosis may also add relevant information for the portfolio choice problem, especially for very fat-tailed assets. This extension to our investor's problem goes beyond the main analysis in the current paper, and is left for future work.

As in Harvey *et al.* (2010), the parameters φ_v and φ_s in the objective function of Equation (19.3) can take arbitrary values to describe particular investor's preferences. Alternatively, following Guidolin and Timmermann (2008) and Jondeau and Rockinger (2012), among others, the three-moment preference assumption could be interpreted as the expected value of the third-order Taylor series expansion of a power utility function with coefficient of risk aversion $\mathcal{A}$. That is,

$$\begin{aligned} \mathbb{E}[U(1 + R_{t+1}(\omega_t))] &\approx \mathbb{E}(R_{t+1}(\omega_t)) - \frac{\mathcal{A}}{2} \mathbb{E}(R_{t+1}(\omega_t)^2) \\ &\quad + \frac{\mathcal{A}(\mathcal{A}+1)}{6} \mathbb{E}(R_{t+1}(\omega_t)^3). \end{aligned} \quad (19.4)$$

where $U(W) = \frac{W^{1-\mathcal{A}}}{1-\mathcal{A}}$ for $\mathcal{A} > 1$ and $U(W) = \log(W)$ for $\mathcal{A} = 1$. In this case, the impact of variance and skewness in the investor's decision rule would be a function of the coefficient of risk aversion $\mathcal{A}$ ($\varphi_v = \frac{\mathcal{A}}{2}$ and $\varphi_s = \frac{\mathcal{A}(\mathcal{A}+1)}{6}$).

Finally, in Equation (19.3), the domain $\mathcal{D} \subset \mathbb{R}^N$ represents the budget constraint defined by

$$\mathcal{D} = \left\{ (\omega_t^1, \dots, \omega_t^N) : \sum_{j=1}^{N-n} \omega_t^j + \sum_{i=N-n+1}^N |\omega_t^i| = 1 \right\}. \quad (19.5)$$

Because both long and short positions in commodity futures contracts require the same initial collateral, we have to take the absolute value of the futures weights $(\omega_t^{N-n+1}, \dots, \omega_t^N)$, such that short positions in futures contracts cannot be used to increase holdings of other assets. Furthermore, if the investor is short-sales constrained, the weights of spot contracts must satisfy that $(\omega_t^1, \dots, \omega_t^{N-n}) \in (0, 1)^{N-n}$.

Once we have the set-up of the investor's problem, we need to define density forecasts of the return joint distribution in order to compute the optimal portfolio weights. In the next section, we describe our multivariate density model.

19.3 Multivariate Conditional Copula with Asymmetry

We employ multivariate conditional copulas to obtain a flexible model for the multivariate distribution of assets' log-returns $\mathbf{r}_{t+1} = (r_{1,t+1}, \dots, r_{d,t+1})$, where $d \leq N$ is the number of risky assets. Every d -variate distribution consists of d marginal distribution functions or *margins* that describe each univariate return's behaviour, as well as a joint dependence function that defines the relations among individual processes. Unlike traditional multivariate distributions, such as the Gaussian and Student's t distributions, copula models support the construction of multivariate distributions with arbitrary univariate processes and dependence.

Formally, a d -variate copula is a d -dimensional distribution function on the unit interval $[0, 1]^d$, that is, a joint distribution with d uniform marginal distributions. Consider a multivariate conditional distribution $F_t(r_{1,t+1}, \dots, r_{d,t+1})$ formed by d univariate conditional distributions $F_{i,t}(r_{i,t+1})$, where the

subscript t denotes that the joint and marginal distributions are conditioned on the information set $\mathcal{F}_t$ available at time t . Following Patton (2006b), there must exist a function C_t that maps the domain $[0, 1]^d$ toward the interval $[0, 1]$, called the *conditional copula*, such that

$$F_t(r_{1,t+1}, \dots, r_{d,t+1}) = C_t(F_{1,t}(r_{1,t+1}), \dots, F_{d,t}(r_{d,t+1})). \quad (19.6)$$

This expression constitutes a d -dimensional extension of Sklar's (1959) theorem for conditional copulas.*

Using the expression in Equation (19.6), any copula C_t can be employed to define a joint distribution $F_t(r_{t+1})$ with the arbitrary marginal distributions $F_{1,t}, \dots, F_{d,t}$. Thus, using a bottom-up approach, we model the marginal distributions of asset returns, followed by the conditional copula function that describes their dependence structure.

19.3.1 Modelling Univariate Processes

We first specify the univariate distribution functions of the asset returns r_{t+1} . Our multivariate copula model supports the use of various marginal distributions. Thus, we can attend to the particular characteristics of each asset return, which is an useful feature when different types of assets appear in the portfolio, such as commodities and stocks. We present a marginal distribution model that captures individual skewness and heavy tails, as well as time-varying moments. We build on the autoregressive conditional density models of Hansen (1994), Harvey and Siddique (1999) and Jondeau and Rockinger (2003), and we propose a generalized Student's t distribution with possibly time-varying parameters. Thus, the univariate process for each asset return $r_{i,t+1}$ ($i = 1, \dots, d$) can be expressed as follows:

$$r_{i,t+1} = \mu_{i,t+1} + \sqrt{\sigma_{i,t+1}^2} z_{i,t+1} \quad (19.7)$$

$$\mu_{i,t+1} \equiv \mathbb{E}_t(r_{i,t+1}) = \mu_{0,i} + \beta'_i X_t + \sum_{j=1}^p \Phi_{i,j} r_{i,t+1-j} \quad (19.8)$$

$$\begin{aligned} \sigma_{i,t+1}^2 \equiv \text{Var}_t(r_{i,t+1}) = & \alpha_{0,i} + \alpha_{1,i}^+ \sigma_{i,t}^2 z_{i,t}^2 1\{z_{i,t} \geq 0\} \\ & + \alpha_{1,i}^- \sigma_{i,t}^2 z_{i,t}^2 1\{z_{i,t} < 0\} + \alpha_{2,i} \sigma_{i,t}^2, \end{aligned} \quad (19.9)$$

$$z_{i,t+1} \sim g(z_{i,t+1}; \nu_{i,t+1}, \lambda_{i,t+1}), \quad (19.10)$$

where $\mu_{i,t+1}$ and $\sigma_{i,t+1}^2$ are the mean and variance conditioned on $\mathcal{F}_t$ for the i th asset returns, and $z_{i,t+1}$ is the corresponding residual.

The conditional mean, defined in Equation (19.8), is a linear function of p lagged returns, $r_{i,t+1-j}$ ($j = 1, \dots, p$), and m further explanatory variables X_t , with the coefficients $\Phi_{i,j}$ and β_i , respectively, and the drift parameter $\mu_{0,i}$. This specification can capture the possible presence of autocorrelation and predictability in asset returns. As the exogenous regressors X_t we consider explanatory variables employed in previous literature to predict variation in stocks and commodity futures returns, including the short rate, default spread, momentum, basis and growth in open interest (see Hong and Yogo (2012)).

As we describe in Equation (19.9), we employ an asymmetric or leveraged GARCH dynamic for the conditional variance. This specification is designed to account for volatility clustering and leverage effects, such as possible asymmetric responses to positive and negative shocks that have occurred

* This generalization of Sklar's theorem is a direct application of the concept of a conditional copula (Patton 2006b, Theorem 1) to a multivariate case (Nelsen 2006, Theorem 2.10.9), and requires simply that conditioning variables be the same for all marginal distributions and the copula. If margins are continuous, this copula is unique.

in the previous period (Campbell and Hentschel 1992).^{*} Equation (19.10) then denotes that the univariate innovations $z_{i,t+1}$ are drawn from a generalized Student's t distribution g , which can capture heavy tails and individual skewness through the degrees of freedom v_i and asymmetry parameter λ_i (Hansen 1994).[†]

Finally, our specification of the marginal distributions addresses the possibility of time-varying higher moments as follows:

$$v_{i,t+1} = \Lambda_{(2,\infty)} \left(\delta_{0,i} + \delta_{1,i}^+ z_{i,t} \mathbf{1}\{z_{i,t} \geq 0\} + \delta_{1,i}^- z_{i,t} \mathbf{1}\{z_{i,t} < 0\} + \delta_{2,i} \Lambda_{(2,\infty)}^{-1}(v_{i,t}) \right), \quad (19.11)$$

$$\lambda_{i,t+1} = \Lambda_{(-1,1)} \left(\zeta_{0,i} + \zeta_{1,i}^+ z_{i,t} \mathbf{1}\{z_{i,t} \geq 0\} + \zeta_{1,i}^- z_{i,t} \mathbf{1}\{z_{i,t} < 0\} + \zeta_{2,i} \Lambda_{(-1,1)}^{-1}(\lambda_{i,t}) \right), \quad (19.12)$$

where $\delta_{0,i}, \delta_{1,i}^+, \delta_{1,i}^-, \delta_{2,i}, \zeta_{0,i}, \zeta_{1,i}^+, \zeta_{1,i}^-$ and $\zeta_{2,i}$ are constant parameters, and $y \equiv \Lambda_{(l,u)}(x) = (u + le^x)/(1 + e^x)$ denotes the modified logistic map designed to keep the transformed variable y in the domain (l, u) for all $x \in \mathbb{R}$. Thus, shape parameters $v_{i,t+1}$ and $\lambda_{i,t+1}$ may depend on their lagged values and react differently to positive and negative shocks. This general specification also includes some well-known univariate distributions as particular cases. For instance, if the asymmetry parameter is 0, we obtain the symmetric Student's t distribution; as degrees of freedom tend to infinity, it converges to a Gaussian distribution.

19.3.2 Modelling Copula Functions

In this section, we present the copula functions that determine the dependence structure of our model. Following Sklar's theorem in Equation (19.6), the copula function acts like a joint distribution of the probability transformed vector $(F_{1,i}(r_{1,t+1}), \dots, F_{d,i}(r_{d,t+1}))'$, where $F_{i,t}(r_{i,t+1})$ are the marginal distribution functions of asset returns $r_{i,t+1}$, as described in Equations (19.7)–(19.10). In particular, we employ three multivariate copula functions: two well-known elliptical copulas, the Gaussian and the t copula (Embrechts *et al.* 2003), and an asymmetric multivariate dependence, the so-called skewed t copula (Demarta and McNeil 2005). They are all implicit dependence functions of various multivariate normal mixtures. More specifically, they are the parametric copula functions contained in the multivariate Gaussian, Student's t and generalized hyperbolic skewed t distributions, respectively.

Through a direct application of Sklar's theorem, we can obtain these implicit copulas by evaluating a given multivariate distribution (e.g. generalized hyperbolic skewed t) at the quantile functions of its corresponding marginal distributions. For instance, the skewed t copula is given by:

$$C^{\text{SK}}(u_1, \dots, u_d; \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma}) = H(H_1^{-1}(u_1; \boldsymbol{\eta}, \boldsymbol{\gamma}_1), \dots, H_d^{-1}(u_d; \boldsymbol{\eta}, \boldsymbol{\gamma}_d); \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma}), \quad (19.13)$$

where $H(\cdot; \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma})$ is the generalized hyperbolic skewed t distribution with $d \times d$ correlation matrix $\mathbf{P}$, degrees of freedom $\boldsymbol{\eta}$, and d -dimensional asymmetry parameter vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_d)$. The $H_i(\cdot; \boldsymbol{\eta}, \boldsymbol{\gamma}_i)$ are the d univariate skewed t distributions, the H_i^{-1} are the corresponding quantile functions, and $(u_1, \dots, u_d)'$ is the probability-transformed vector. Similarly, we can extract the Gaussian and t copulas, $C^G(u_1, \dots, u_d; \mathbf{P})$ and $C^T(u_1, \dots, u_d; \mathbf{P}, \boldsymbol{\eta})$, from their respective multivariate distributions. In Appendix A, we derive explicitly their density functions.

Using these three copulas, we can model three different types of dependence. The Gaussian copula, $C^G(\cdot; \mathbf{P})$, defines linear, symmetric dependence, completely determined by the correlation matrix $\mathbf{P}$. Thus, it is unable to capture tail dependence or asymmetries. The t copula, $C^T(\cdot; \mathbf{P}, \boldsymbol{\eta})$, is also elliptically symmetric but allows for tail dependence through the degrees of freedom parameter, $\boldsymbol{\eta}$. The t copula

^{*} To guarantee positive and stationary volatility, the parameters of the variance dynamics in Equation (19.9) must satisfy the following constraints: $\alpha_{0,i} > 0$; $\alpha_{1,i}^+, \alpha_{1,i}^-, \alpha_{2,i} \geq 0$ and $\alpha_{2,i} + (\alpha_{1,i}^+ + \alpha_{1,i}^-) / 2 < 1$.

[†] The technical details of this univariate distribution can also be found in Jondeau and Rockinger (2003). We summarize them briefly in the supplemental data.

assigns more probability to the extremes than does the Gaussian copula. The greater the degrees of freedom, the smaller the level of tail dependence, converging in the limit as $\eta \rightarrow \infty$ to the Gaussian copula. Finally, the skewed t copula, $C^{\text{SK}}(\cdot; \mathbf{P}, \eta, \boldsymbol{\gamma})$, can capture extreme and asymmetric dependence of the asset returns. Through the d -dimensional vector of asymmetry parameters $\boldsymbol{\gamma}$, the skewed t copula can assign more weight to one tail than the other. For example, if all elements of the asymmetry vector are negative, the density contour of the skewed t copula would be clustered in the negative-negative quadrant. Eventually, if $\boldsymbol{\gamma} \rightarrow \mathbf{0}$, asymmetric dependence goes to 0, and we would recover the symmetric t copula.

Once we have defined the functional forms of the three implicit copulas, we can build the multivariate distribution model for our vector of asset returns. This multivariate distribution is formed from the marginal distributions of the previous section and one of the implicit copulas we described previously. In addition, following pioneering works by Patton (2006a,b), we can parameterize time variation in the conditional copula function of our multivariate model. For that purpose, and in the spirit of Engle's (2002) dynamic conditional correlation model, we extend the notion to other types of dependence beyond the Gaussian one and allow that the dependence matrix $\mathbf{P}_t$ of our conditional copula may evolve over time, according to some GARCH-type process. In the most general case, the vector of return innovations, $\mathbf{z}_{t+1} = (z_{1,t+1}, \dots, z_{d,t+1})'$, follows the distribution specification:

$$\mathbf{z}_{t+1} \sim C^{\text{SK}}(g(z_{1,t+1}; \nu_{1,t+1}, \lambda_{1,t+1}), \dots, g(z_{d,t+1}; \nu_{d,t+1}, \lambda_{d,t+1}); \mathbf{P}_{t+1}, \eta, \boldsymbol{\gamma}), \quad (19.14)$$

where the $g(z_{i,t+1}; \nu_{i,t+1}, \lambda_{i,t+1})$ are the conditional univariate distributions in Equation (19.10), and the evolution equation for $\mathbf{P}_{t+1}$ is given by:

$$\mathbf{P}_{t+1} = \Lambda_{(-1,1)} \left(\kappa_0 \bar{\mathbf{P}} + \kappa_1 \frac{1}{M} \sum_{m=1}^M x_{t+1-m} x'_{t+1-m} + \kappa_2 \mathbf{P}_t \right). \quad (19.15)$$

In this case, x_t is the vector of transformed variables $(H^{-1}(u_{1t}; \eta, \gamma_1), \dots, H^{-1}(u_{dt}; \eta, \gamma_d))'$; $\bar{\mathbf{P}}$ is the constant correlation matrix; κ_0 , κ_1 and κ_2 are constant parameters; and M is the number of lags we consider. The modified logistic function $\Lambda_{(-1,1)}(\cdot)$ ensures that the elements of $\mathbf{P}_{t+1}$ remain in the domain $(-1, 1)$.

Using multivariate copulas to model high dimensional dependence is still an open problem. An advantage of using our parametric approach is that, a priori, it can be generalized to a finite number of dimensions. In addition, our model can be interpreted in terms of other more traditional, well-known dependence models, which are included in our specification as nested cases. One of the costs of using these parametric models for the joint distribution is that the number of parameters (i.e. correlation coefficients and asymmetry vector) increase significantly with the number of assets in the portfolio. In our three-dimensional problem for oil, gold and equity returns, the number of parameters in the model is similar to other conditional copula models proposed in the literature (Patton 2004, 2006b, Jondeau and Rockinger 2006). In the empirical application, to avoid in-sample overfitting and spurious findings, we carry out a detailed out-of-sample analysis. Furthermore, we also consider other nested, more parsimonious, specifications, as benchmarks to check the consistency of our estimators.

Other recent works, such as Härdle *et al.* (2010), have proposed other interesting procedures, based on generalizations of Archimedean copulas, to deal with time-varying, non-Gaussian dependence, but using a small number of parameters. See also Manner and Reznikova (2012) for a survey on the growing field of time-varying copulas.

19.3.3 Model Estimation and Portfolio Optimization

Our model structure, formed by the marginal distributions and the copula, allows for a two-step estimation procedure, similar to the conditional set-ups of Jondeau and Rockinger (2006) and Patton (2006a). In the first step, we obtain the maximum likelihood (ML) estimates of the individual return processes;

then, we determine the parameter estimates of the copula function.* From this ML approach, we can compute the asymptotic and robust standard errors for the estimates.

Formally, this procedure can be expressed as follows: Let $\bar{\mathbf{r}}_T = \{\mathbf{r}_1, \dots, \mathbf{r}_T\}$ be the sample of returns of length t , where $\mathbf{r}_t = (r_{1,t}, \dots, r_{d,t})$ for $t = 1, \dots, T$. We want to find the set of parameter estimates $\hat{\boldsymbol{\theta}}$ that maximizes the log-likelihood function $\mathcal{L}$, that is,

$$\hat{\boldsymbol{\theta}} \equiv \arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta}; \bar{\mathbf{r}}_T) = \arg \max_{\boldsymbol{\theta} \in \Theta} \sum_{t=1}^T \log f_t(r_{t+1}; \boldsymbol{\theta}), \quad (19.16)$$

where $f_t(r_{t+1}; \boldsymbol{\theta})$ is the probability density function of the multivariate model conditioned by the information set $\mathcal{F}_t$ and parameterized by $\boldsymbol{\theta} \in \Theta$.

From the assumptions of Sklar's theorem in Equation (19.6), we can decompose the log-likelihood function $\mathcal{L}$ in Equation (19.16) into two parts, the margins and the copula (see the details in [Appendix B](#)):

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_M, \boldsymbol{\theta}_C; \bar{\mathbf{r}}_T) &= \sum_{i=1}^d \mathcal{L}_i(\boldsymbol{\theta}_{i,M}; \bar{\mathbf{r}}_T) + \mathcal{L}_C(\boldsymbol{\theta}_C; \boldsymbol{\theta}_M, \bar{\mathbf{r}}_T) \\ &= \sum_{i=1}^d \sum_{t=1}^T \log f_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_{i,M}) + \sum_{t=1}^T \log c_t(u_{1,t+1}, \dots, u_{d,t+1}; \boldsymbol{\theta}_C). \end{aligned} \quad (19.17)$$

where $\mathcal{L}_i$ and $\mathcal{L}_C$ are the log-likelihood functions for the i th marginal process and the copula, $\boldsymbol{\theta}_M$ denotes the set of parameters corresponding to the d marginal distributions, $(\boldsymbol{\theta}_{1,M}, \dots, \boldsymbol{\theta}_{d,M})'$, and $\boldsymbol{\theta}_C$ denotes the parameters of the copula function. The $u_{i,t+1}$ are the marginal distributions $F_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_{i,M})$, with corresponding marginal density functions $f_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_{i,M})$; and $c_t(\cdot; \boldsymbol{\theta}_C)$ is the copula density function.

Once we have estimated the model density function, we use this information to obtain the optimal portfolio. For our parametric density models, the integrals defining the portfolio return moments involved in the investor's optimization problem of Equation (19.3) do not have a closed-form solution. Using Monte Carlo simulations to estimate the value of these integrals, we can solve numerically the optimization problem. In this respect, an advantage of our implicit copulas is that it is easy to sample from them, as long as we are able to sample from the normal mixture distribution from which they are extracted. Similar approaches to simulate heavy-tailed data and copulas in the context of portfolio optimization have been implemented by Patton (2004), Biglova *et al.* (2009) and Ortobelli *et al.* (2010) (and the references therein). The following procedure outlines the implementation of the model estimation and the portfolio optimization:

1. Following Equation (19.17), estimate sequentially the d marginal distributions models and the copula function:†

$$\hat{\boldsymbol{\theta}}_{i,M} = \arg \max_{\boldsymbol{\theta}_{i,M}} \mathcal{L}_i(\boldsymbol{\theta}_{i,M}; \bar{\mathbf{r}}_T), \text{ for } i = 1, \dots, d, \text{ and } \hat{\boldsymbol{\theta}}_C = \arg \max_{\boldsymbol{\theta}_C} \mathcal{L}_C(\boldsymbol{\theta}_C; \hat{\boldsymbol{\theta}}_M, \bar{\mathbf{r}}_T).$$

In our model specification, the set of parameters for margin i is given by $\boldsymbol{\theta}_{i,M} = (\mu_{0,i}, \boldsymbol{\beta}'_i, \Phi_{i,1}, \Phi_{i,2}, \Phi_{i,3}, \alpha_{0,i}, \alpha_{1,i}^+, \alpha_{1,i}^-, \alpha_{2,i}, \delta_{0,i}, \delta_{1,i}^+, \delta_{1,i}^-, \delta_{2,i}, \zeta_{0,i}, \zeta_{1,i}^+, \zeta_{1,i}^-, \zeta_{2,i})'$ (see [Equations \(19.8\)](#), [\(19.9\)](#), [\(19.11\)](#), and [\(19.12\)](#)), and the set of parameters for the skewed t copula is given by $\boldsymbol{\theta}_C = (\mathbf{P}_C, \kappa_0, \kappa_1, \kappa_2, \eta, \gamma)'$ (see [Equations \(19.14\)](#) and [\(19.15\)](#)).

* This procedure is also known as the inference functions for margins method. A similar two-stage approach is used to estimate some multivariate GARCH models, such as the constant (CCC) and dynamic (DCC) conditional correlation models (see Engle and Sheppard (2001)).

† Patton (2006a) shows that one-step maximum likelihood estimators and two-stage estimators are equally asymptotically efficient.

Some remarks should be considered though. First, the quality of the copula estimation depends strongly on the goodness of fit of the parametric functions we use for the marginal distribution models. Second, t copulas require the estimation of shape parameters η and γ , apart from the correlation matrix $\mathbf{P}$; in these cases, because optimization of the objective function often falls into local maxima, convergence difficulties may arise when maximizing the log-likelihood function $\mathcal{L}_C$ directly. To overcome this problem, we perform the estimation of the t copulas using an iterative procedure: An inner function computes $\hat{\mathbf{P}}_t$, maximizing the likelihood given the values of the shape parameters; then, this function is maximized with respect to the shape parameters.*

2. Use the parameter estimates $\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\theta}}_M, \hat{\boldsymbol{\theta}}_C)'$ and the dynamics of the return vector $\mathbf{r}_{t+1}$, defined in Equations (19.7)–(19.10) and (19.14), to obtain the forecast density for the next period; then, generate $Q = 10,000$ independent draws $\{\mathbf{r}_t^q\}_{q=1}^Q$ sampling from that density:
 - a. Generate Q random vectors $\{\mathbf{u}_{t+1}^q\}_{q=1}^Q$ from the implicit copulas. For example, for the skewed t copula: $\mathbf{u}_{t+1}^q = (H_1(x_1^q; \hat{\eta}, \hat{\gamma}_1), \dots, H_d(x_d^q; \hat{\eta}, \hat{\gamma}_d))$, where the $(x_1^q, \dots, x_d^q)$ are random vectors from the multivariate distribution $H(\mathbf{0}, \hat{\mathbf{P}}_t, \hat{\eta}, \hat{\gamma})$ generated using Equation (A.1).† Similarly, we can sample from the Gaussian and t copulas.
 - b. Use the inverse functions of the conditional marginal distributions to obtain Q draws of the innovations $\{z_t^q\}_{q=1}^Q$ where $z_t^q = (g^{-1}(u_{1,t+1}^q; \hat{\nu}_{1,t}, \hat{\lambda}_{1,t}), \dots, g^{-1}(u_{d,t+1}^q; \hat{\nu}_{d,t}, \hat{\lambda}_{d,t}))$; then employ the forecast conditional mean and variance to generate $\{\mathbf{r}_t^q\}_{q=1}^Q$.
3. To obtain the optimal portfolio weights $\boldsymbol{\omega}^*$, maximize the investor's objective function in Equation (19.3) subject to the non-linear budget constraint of Equation (19.5) using the generated simulations to estimate the moments of the portfolio return.

19.4 Empirical Application

In this section, we first present the data and their main univariate and multivariate statistical properties. Then, we estimate the conditional copula models and analyse their in-sample fitting performance. Finally, we solve the portfolio problem for the copula models numerically and obtain the optimal weights, investment ratios and relative performance measures over the out-of-sample period.

19.4.1 Data and Preliminary Analysis

Our empirical application relies on three risky assets: two commodity futures, oil and gold and the S&P 500 equity index. The oil futures correspond to West Texas Intermediate (WTI) crude oil from the New York Mercantile Exchange (NYMEX). The gold futures correspond to the gold bar, with a minimum of 0.995 fineness, from the New York Commodities Exchange (COMEX). These futures are two of the most actively traded commodity contracts in the world, and they do not have tight restrictions on the size of daily price movements.‡ In both cases, we employ the most liquid futures contracts, measured by daily trading volume, of all maturities available. The risk-free rate is computed from the three-month US Treasury bills provided by the Federal Reserve System. All data are in US dollars and came from Thomson-Reuters Datastream. The sample period considered ranges from 20 June 1990 to 8 September 2010, for a total of 1056 weekly observations. We divided the sample in two subperiods, such that the period from 20 June 1990 to 20 June 2006 (836 observations) supported the in-sample estimation

* Furthermore, we employ a global optimization approach, consisting of simulated annealing (Goffe *et al.* 1994), to check the validity of the local optimization results.

† Note that, for the estimation and simulation of the skewed t copula, to improve the feasibility of the computations related to the cumulative density and inverse functions of the univariate generalized hyperbolic skewed t distribution, H_i and H_i^{-1} , we approximate its density function h_i using cubic splines.

‡ At the end of 2011, gold and crude oil futures represented 30% of the Dow Jones-UBS Commodity Index and 38% of the S&P-Goldman Sachs Commodity Index.

analyses of the models, and the remaining 220 observations from 20 June 2006 to 8 September 2010 were reserved for the out-of-sample portfolio performance exercise.

19.4.1.1 Univariate Analysis

We first analyse the univariate behaviour of the three asset returns. In the supplemental data, we report the summary statistics for the weekly returns of the gold and oil futures and the equity index for the sample periods, as well as plots of the relative price moves of each asset over the full-sample period. Here, we summarize the main findings.

We observe substantial changes in the sample moments of returns over time. The mean returns are all positive, except for equity during the out-of-sample period (June 2006–September 2010). Return volatilities per week for oil, gold and equity increased from 4.4, 1.9 and 2.1% in the in-sample period to 5.4, 3.2 and 3.0% during the out-of-sample period. Looking at the ratio of the mean over the volatility (Sharpe's ratio), we find that for the in-sample observations, equity (0.07) performs better than oil (0.04) and gold (0.03). This pattern changed during the 2006–2010 period, during which ratios of oil (0.01) and equity (−0.02) were below their historical average, whereas gold's ratio (0.11) moved significantly above its historical average. According to the Ljung-Box (LB) and Lagrange multiplier (LM) statistics, there is evidence of serial correlation in the returns and squared returns for all time series (except for oil returns over the in-sample period).

Assets returns are non-normal, skewed and heavy tailed. According to the Jarque-Bera (JB) and Kolmogorov-Smirnov (KS) tests, normality in the returns' unconditional distribution is strongly rejected for all samples. Skewness and kurtosis of returns differs across assets and sample periods. From the in-sample to the out-of-sample period, the equity returns' skewness grew much more negative, while gold returns changed from positive to negative skewness, and oil returns from negative to positive. During 2006–2010, the oil and gold returns' kurtosis decreased with respect to the previous period, but equity returns' kurtosis strongly increased, as expected.

19.4.1.2 Multivariate Analysis

It is also important to describe the interactions observed in the sample among the oil, gold and equity index returns. We first focus on the characteristics of the linear dependence; then we turn to analysing non-linear features observed in the vector of returns. Details of the multivariate statistics and tests for the three-dimensional vector of asset returns are reported in the supplemental data.

We find a large increase in linear dependence for the 2006–2010 period with respect to historical values. The sample correlation between oil and gold returns rises from 0.21 to 0.39. Furthermore, equity index returns, which were negatively correlated with oil (−0.06) and gold (−0.08) in 1990–2006, became positively correlated with both of these commodity returns over the 2006–2010 period, with coefficients equal to 0.39 and 0.17, respectively. These findings suggest that dependence between commodities and equity is no longer constant and evolves with time. To check this assumption, we carry out Engle and Sheppard's (2001) test for constant correlation. The probabilities of constant correlation (test p -values) are less than 0.05 in all cases; therefore, we reject the hypothesis of constant dependence.

We calculate the trivariate measures of skewness and kurtosis proposed by Mardia (1970) to test multivariate normality. The corresponding statistics suggest that the hypothesis of multivariate normality should be rejected for the two sample periods considered. In addition, following McNeil *et al.* (2005), we test whether the standardized vector of returns is consistent with a spherical distribution. The corresponding KS test statistics reject the ellipticity hypothesis for all samples. Visually, their associated quantile-quantile plots reveal that multivariate normality and elliptical symmetry are strongly rejected for our sample (see supplemental data).

Finally, to check for the presence of asymmetric dependence between asset returns in our sample, we analysed the exceedance correlation and tail dependence. For each pair of asset returns, we thus, compute and plot the exceedance correlation function proposed in Ang and Chen (2002) and Longin and Solnik (2001), which depicts the correlation between returns above or below a given quantile. In the

case of symmetric dependence, the correlation for both extremes should be similar and equal to zero for Gaussian dependence. According to these plots, any assumptions of normality or symmetry seem unrealistic for our sample. Oil and gold do not display the same level of diversification for bear and bull markets, and correlation between oil and equity is highly positive for large negative returns but smaller for large positive returns. The correlation between gold and equity is close to 0 for large negative returns and significantly positive for very large positive returns. Although oil and gold are very positively correlated for large negative returns, are not, or even are negatively, correlated for large positive returns. Büyüksahin *et al.* (2010) also find patterns of extreme and asymmetric dependence between commodity indexes and equity.

Furthermore, when we estimate the tail dependence of each pair of returns in our sample, we also observe an asymmetric pattern. In particular, we fit the upper and lower tail dependence parameters, τ^U and τ^D , corresponding to the symmetrized Joe-Clayton (SJC) copula, defined by Patton (2006b). We observe that tail dependence increases over the 2006–2010 period, and lower tail dependence estimates are generally larger than the upper ones, especially in the latter sample period (see supplemental data for details).

In summary, both univariate and multivariate analyses suggest that the assumptions of normality and symmetry for the individual processes and dependence functions are very restrictive and should probably be rejected. A flexible model that captures all the features analysed in the data is thus required. In the next section, we estimate the conditional copula model proposed in Section 19.3 for our vector of oil, gold and equity returns. Subsequently, we investigate whether capturing these features (e.g. non-normality of the individual processes, time-varying moments and asymmetric dependence) using the more flexible model leads to economically better portfolio decisions.

19.4.2 Estimation of the Conditional Copula Model

In this section, we estimate the conditional copula model using the multistage maximum likelihood procedure explained in Section 19.3. We first present the in-sample estimation and goodness-of-fit test results for the marginal distribution models. In a second stage, we analyse the results for the copula model.

19.4.2.1 In-Sample Results for the Marginal Distributions

Table 19.1 presents maximum likelihood estimates of the parameters of the marginal distributions for oil, gold and equity index returns. We compute robust standard errors of these estimates and report their corresponding p -values in parentheses. These estimates correspond to the generalized t marginal distribution function with time-varying moments, described in Equations (19.7)–(19.12) of Section 19.3.

In the mean equation, we find that the basis (for oil returns) and momentum and risk-free rate (for gold returns) are significant explanatory factors. The results for the variance equation further show that volatility is strongly persistent for all returns. For equity, only negative returns have an effect on subsequent variance. This result is consistent with the leverage effect studied by Campbell and Hentschel (1992), among others. Yet for both commodities, especially gold, we observe an *inverse* leverage effect; that is, positive shocks have a stronger effect on variance than do negative ones of the same size.

Regarding the dynamics of the degrees-of-freedom and asymmetry parameters, we find that both higher moments are rather persistent for all asset returns over the in-sample period. Large moves in oil returns, especially negative ones, diminish the posterior degrees of freedom ($\delta_1^- = 5.51 (0.01)$, $\delta_1^+ = -0.73 (0.58)$, with p -values in parentheses), increasing the likelihood of posterior large shocks. For equity returns, large moves, especially positive ones, increase the subsequent degrees of freedom ($\delta_1^- = -2.49 (0.38)$, $\delta_1^+ = 21.14 (0.05)$), so large returns are less likely. For gold returns, extreme events are generally more likely to cluster in periods of large positive moves: Positive shocks are followed by a decrease in posterior degrees of freedom ($\delta_1^+ = -3.63 (0.00)$), whereas negative shocks generally are followed by an increase ($\delta_1^- = -3.26 (0.00)$).

TABLE 19.1 In-Sample Results for the Marginal Distribution Models.

	Oil		Gold		Equity	
	<i>coeff.</i>	<i>(p-val.)</i>	<i>coeff.</i>	<i>(p-val.)</i>	<i>coeff.</i>	<i>(p-val.)</i>
Mean equation						
μ (/100)	0.140	(0.316)	0.474	(0.037)	0.130	(0.037)
$basis_{t-1}$	-0.031	(0.216)				
$momentum_{t-1}$			-0.175	(0.153)		
r_{t-1}^f			-0.111	(0.013)		
r_{t-1}					-0.108	(0.002)
r_{t-2}			-0.041	(0.213)		
r_{t-3}	0.050	(0.110)				
Variance equation						
α_0 (/1000)	0.025	(0.175)	0.013	(0.012)	0.016	(0.022)
a_1^+	0.085	(0.001)	0.183	(0.001)	0.000	(0.956)
a_1^-	0.060	(0.030)	0.029	(0.189)	0.145	(0.000)
α_2^1	0.920	(0.000)	0.870	(0.000)	0.890	(0.000)
Degrees-of-freedom equation						
δ_0	0.100	(0.000)	0.025	(0.042)	-0.200	(0.054)
δ_1^+	-0.732	(0.576)	-3.633	(0.000)	21.138	(0.049)
δ_1^-	5.514	(0.007)	-3.257	(0.002)	-2.489	(0.383)
δ_2	0.998	(0.000)	1.009	(0.000)	0.966	(0.000)
Asymmetry parameter equation						
ζ_0 (/10)	0.085	(0.177)	0.088	(0.111)	-0.250	(0.099)
ζ_1^+	-0.357	(0.313)	-0.903	(0.041)	-0.374	(0.798)
ζ_1^-	0.218	(0.702)	0.592	(0.222)	-2.553	(0.034)
ζ_2	0.998	(0.000)	1.001	(0.000)	0.981	(0.000)
$logL$	1515.4		2262.7		2149.6	
$\bar{v}$	11.24	(0.002)	4.792	(0.000)	12.25	(0.009)
$\bar{\lambda}$	-0.093	(0.070)	0.018	(0.577)	-0.230	(0.000)
$logL_{const}$	1503.2		2253.0		2140.5	

Notes: This table reports the maximum likelihood parameter estimates of the marginal distribution model for oil, gold and equity-index returns with generalized Student's t distribution and time-varying moments. Parameters of the mean, variance, degrees of freedom and asymmetry are defined in Equations (19.7), (19.9), (19.11) and (19.12), respectively. The results correspond to the estimation period from June 1990 to June 2006 (836 observations). The p -values of the estimates appear in parentheses and are computed using the robust standard errors. $logL$ is the log-likelihood of the marginal distribution model. $\bar{v}$, $\bar{\lambda}$ and $logL_{const}$ are respectively the degrees of freedom, asymmetry parameter and log-likelihood of the constant version of the model.

In general, lagged values of the asymmetry parameter are more significant for subsequent parameter values than is the effect of the previous returns shock. Over our study's in-sample period, only positive shocks in gold returns ($\zeta_1^+ = -0.90(0.04)$) and negative shocks in equity returns ($\zeta_1^- = -2.55(0.03)$) seem to have effects of opposite sign on the posterior asymmetry parameters. Therefore, for the three assets returns, we find significant time variation in the moments of the univariate processes. As a benchmark, we also estimate the degrees-of-freedom and asymmetry parameters, $\bar{v}$ and $\bar{\lambda}$, of the conditional distribution with constant shape parameters. We find that for the in-sample period, the left tail of the conditional distribution of oil and equity returns is fatter than the right tail, with parameters $\bar{v}$ and $\bar{\lambda}$ equal to 11.24 and -0.09 for oil returns, and 12.25 and -0.23 for equity returns. In contrast, gold returns have positive (though not significant) asymmetry parameter ($\bar{\lambda} = 0.02$) and heavier tails than oil and equity returns ($\bar{v} = 4.79$).

A reliable estimation of copula models requires an appropriate specification of the univariate density functions (see Patton (2006a,b) and Jondeau and Rockinger (2006)). Therefore, to avoid misspecified copula models, we conduct the in-sample goodness-of-fit test suggested by Diebold *et al.* (1998) for our estimated marginal distribution models (results are reported in the supplemental data). If the marginal model is correctly specified, the probability integral transform should be i.i.d. Uniform(0,1). According to the ML statistics, we must reject serial dependence in the first four moments of the probability integral transform (all p -values > 0.15). In addition, the KS statistics suggest that the shape of the conditional distribution model is correctly specified for the three returns (p -values > 0.90). Visually, the goodness-of-fit plots also support these results. Furthermore, the asymmetric marginal model performs substantially better than the Gaussian and symmetric models, even for constant higher moments.

Finally, using likelihood ratio (LR) tests, we compare our more general skewed t marginal model against different constrained alternative models: the generalized t distribution with constant parameters, the standard Student's t distribution with time-varying and constant degrees of freedom, and the standard Gaussian distribution. In all cases, we reject the restricted specification in favour of a more general model, at least at a 5% significance level (p -values < 0.05 ; see detailed results in the supplemental data).

19.4.2.2 In-Sample Results for the Copulas

In this second stage, we estimate the dependence function that links the three marginal distribution models. We analyse the three time-varying or conditional copula functions described in Section 19.3.2: the Gaussian, t , and skewed t copulas. Table 19.2 presents the in-sample ML estimates of the three conditional copula models, estimated on the transformed residuals of the generalized t univariate model. This table also reports the estimated p -values, computed from the asymptotic covariance matrix, and the likelihood values at the optimum for the conditional and unconditional copulas.*

According to the ML estimates, for all copula models, the dependence coefficient between oil and gold returns, $\rho_{oil,gold}$, is positive, whereas the dependence coefficient between gold and equity index returns, $\rho_{gold,equity}$, is negative, and that between oil and equity, $\rho_{oil,equity}$, is insignificantly different from 0. Therefore, the estimated dependence coefficients are consistent with the observed unconditional correlations.

Estimates of the degrees of freedom η are strongly significant for symmetric and skewed t copulas, indicating the presence of a significant level of dependence in the extremes. Regarding the estimation of the skewed t copula, we find that all elements of the asymmetry parameters vector $(\gamma_{oil}, \gamma_{gold}, \gamma_{equity})$ are negative, especially for oil and equity, suggesting more extreme dependence among returns during extreme depreciations of these assets than during bullish markets.

The parameters κ_0 , κ_1 and κ_2 , which parameterized the dynamic equation of dependence, are significant for all conditional copulas, showing strong evidence of time variation and persistence in the conditional dependence. These results regarding the estimates of the dependence functions are also consistent with the preliminary multivariate analysis of Section 19.4.1.2.

According to the LR test statistics reported in Table 19.2, we observe, first, that conditional copulas are preferred over their corresponding unconditional versions (p -values ≤ 0.05 for the three cases). Second, the presence of tail dependence in the in-sample data is not negligible. The p -values of the LR test of the conditional and unconditional t copulas with respect to the more restrictive Gaussian copulas are always less than 0.10. Third, there is evidence of asymmetry dependence over the in-sample period, captured by the skewed t copula, but the gains from modelling this asymmetry may not make up for the penalty associated with the inclusion of more parameters in the model. These gains seem less significant than those obtained from modelling time-varying and tail dependence.

* In this empirical application, the conditional dependence follows the dynamics in equation (15) for $M = 4$, which is the number of lags consistent with the autoregressive lags considered in the univariate models.

TABLE 19.2 In-Sample Results and LR Tests for the Copula Models.

	Conditional copulas					
	Gaussian		<i>t</i> copula		Skewed <i>t</i>	
	<i>coeff.</i>	(<i>p-val.</i>)	<i>coeff.</i>	(<i>p-val.</i>)	<i>coeff.</i>	(<i>p-val.</i>)
$\rho_{oil,gold}$	0.159	(0.000)	0.157	(0.000)	0.161	(0.000)
$\rho_{oil,equity}$	−0.020	(0.556)	−0.016	(0.653)	−0.013	(0.653)
$\rho_{gold,equity}$	−0.064	(0.064)	−0.058	(0.108)	−0.057	(0.108)
η			18.998	(0.025)	19.050	(0.031)
γ_{oil}					−0.268	(0.027)
γ_{gold}					−0.018	(0.934)
γ_{equity}					−0.141	(0.139)
κ_0	0.136	(0.223)	0.128	(0.192)	0.127	(0.265)
κ_1	0.079	(0.055)	0.069	(0.054)	0.060	(0.082)
κ_2	1.647	(0.000)	1.666	(0.000)	1.676	(0.000)
CL		18.680		20.828		23.136
CL _{uncond}		12.453		15.222		18.134
LR(vs. Uncond.)	12.454	(0.006)	11.212	(0.011)	10.004	(0.019)
LR(vs. Gaussian)			4.296	(0.038)	8.913	(0.063)
LR(vs. Symmetric)					4.617	(0.202)
LR(vs. Uncond. Gaussian)			16.750	(0.002)	21.367	(0.003)
LR(vs. Uncond. Symmetric)					15.829	(0.015)
LR(Uncond. vs. Gaussian)			5.538	(0.019)	11.363	(0.023)
LR(Uncond. vs. Symmetric)					5.826	(0.120)

Notes: This table presents the maximum likelihood parameter estimates of the copula models under different assumptions of the conditional joint dependence. The results correspond to the period from June 1990 to June 2006 (836 observations). In each case, the copula is defined by the next set of parameters: the correlation matrix ($\bar{P} = \{\rho_{i,j}\}$), the degrees of freedom (η), the asymmetry vector (γ) and the parameters of the dynamics (κ_0 , κ_1 and κ_2) (see Equation (19.15)). For each parameter estimate, we report in parentheses the *p*-values computed from the asymptotic covariance matrix. The conditional copula likelihood at the optimum is denoted by *CL*, whereas *CL*_{uncond} reports the likelihood at the optimum of the corresponding unconditional version of the copula. We also report the likelihood ratio test statistics (LR) for different restrictive specifications of the copula models. The LR(vs. Uncond.) corresponds to the LR test with respect to the corresponding unconditional version of the copula model. In LR(vs. Symmetric), we test with respect to the conditional *t* copula. In LR(vs. Gaussian), the restrictive model is the conditional Gaussian copula. With LR(vs. Uncond. Symmetric) and LR(vs. Uncond. Gaussian), we test the conditional copulas with respect to the unconditional *t* and Gaussian copulas. Finally, with LR(Uncond. vs. Symmetric) and LR(Uncond. vs. Gaussian), we test the unconditional versions of the copulas with respect to the unconditional *t* and Gaussian copulas.

19.4.2.3 Out-of-Sample Parameter Forecasts

To obtain the optimal portfolio decisions based on our copula models over the out-of-sample period, we need the forecasts of the different parameters at play over the 2006–2010 period. For that purpose, we recursively re-estimate the marginal and copula models throughout the out-of-sample period (220 weekly observations) using a rolling window scheme that drops distant observations as more recent ones are added and therefore keeps the size of the estimation window fixed at 836 observations. Once we re-estimate the model for each point in the out-of-sample period, we construct the time-series of one-period-ahead parameter forecasts needed for the allocation stage (see Section 19.3.3 for the details of the implementation).




FIGURE 19.1 Conditional parameters of the marginal distribution model. Notes: This figure shows the one-step ahead forecasts over the out-of-sample period for the conditional mean, volatility and skewness of the marginal distribution model with generalized Student's t distribution.

Figure 19.1 shows the output of the forecasts of the conditional mean, volatility and skewness of each return process throughout the out-of-sample period. The volatility forecasts of all asset returns are relatively high, especially around October 2008. Conditional skewness is negative for equity and oil returns during the 2006–2010 period, but it is positive for gold returns during that period.

Figure 19.2 presents the forecasts of the conditional dependence parameters. It is worth noting that there is an increase in the fitted correlation coefficients among oil, gold and equity from October 2008, especially for oil and equity returns (see Exhibit 1). In addition, the dependence coefficients seem to evolve more similarly in the latter part of the sample. The degrees-of-freedom forecasts decrease after August 2007, indicating rising tail dependence since then (see Exhibit 2). In addition, the asymmetry parameter of oil ranges between -0.6 and -0.2 , which implies that extreme dependence seems to be stronger during large depreciations of oil, compared with large drops in gold or equity, whose asymmetry parameters range between -0.2 and $+0.2$ (see the forecast of the asymmetry parameter vector in Exhibit 3).

In general, during our re-estimation of the copula models, we find no evidence to contradict skewed and fat-tailed marginal distributions and asymmetric and extreme conditional dependence, but strong evidence indicates that Gaussian distribution and elliptical dependence are not the best-fitting models. These results over the allocation period are consistent with the sample statistics we described previously.

In summary, the skewed t copula provides a more informative measure of the dependence between commodities and equity-index returns, even taking into account that part of the tail behaviour is captured by the skewed fat-tailed marginal distribution models. Therefore, possibly univariate tail behaviour and asymmetric dependence are key factors not taken into account in a standard elliptical, *à la* Markowitz, approach. The extent to which these factors have a significant impact on the portfolio choice decision is addressed in the next section.




FIGURE 19.2 Conditional parameters of the conditional skewed t copula. Notes: This figure shows the one-step ahead forecasts over the out-of-sample period for the correlation coefficients, degrees of freedom and asymmetry vector components of the conditional skewed t copula model.

19.4.3 Optimal Portfolio Results

We now investigate the optimal portfolio decisions based on the copula models we estimated in the previous section, and analyse their performance over the out-of-sample period. In particular, we compare six model-driven portfolio strategies that can be analysed from the perspective of copula models and therefore estimated using the multistage procedure from [Section 19.3.3](#).

First, we consider the unconditional multivariate Gaussian model (Markowitz strategy), a constant Gaussian copula with unconditional Gaussian marginal distributions. Second, we generalize this case by considering two conditional multivariate Gaussian distributions: the constant conditional correlation (CCC) and the dynamic conditional correlation (DCC). Both CCC and DCC specifications are formed by conditional Gaussian marginal distributions with conditional means and variances defined in Equations (19.7) and (19.9). Third, we compute portfolio strategies using the conditional copula models introduced in [Section 19.3](#). Thus, we consider the generalized Student's t distribution for the marginal

models (Equations 19.7–19.10) and three types of conditional dependence functions: the Gaussian, t and skewed t copulas (defined in Equations 19.14 and 19.15). With this set of alternatives, we can compare the gains of including more flexible models as a means to compute portfolio decisions. In addition, we include in the analysis, the equally weighted portfolio, as a common benchmark used in prior literature.

We analyse the portfolio allocations for different parameterizations of the investor's three-moment preferences, defined by φ_v and φ_s . For that purpose, we follow two complementary approaches: one where the specification of the investor's preferences is related to the third-order Taylor series expansion of an utility function with coefficient of relative risk aversion A (as in Guidolin and Timmermann (2008) and Jondeau and Rockinger (2012)), and an other where the different values of φ_v and φ_s account for different arbitrary impacts of the portfolio variance and skewness on the investor's preferences (as in Harvey *et al.* (2010)).

19.4.3.1 Portfolio Weights

In this section, we analyse the time series of portfolio weights over the out-of-sample period corresponding to the portfolio allocations obtained from the six copula models described previously. For that purpose, we report in Table 19.3 the quantiles of the distribution of optimal weights for oil, gold and equity, under various specifications (φ_v and φ_s) of the three-moment investor's preferences, characterized, in this case, by different values of the risk aversion coefficient A . These allocation results are obtained from the unconstrained and short-sales constrained optimizations.* In addition, in Panel A of Figure 19.3 we plot, for two of these preferences specifications, the time-series of portfolio weights resulting from the portfolio decisions made using our most general model, the conditional skewed t copula. Panel B of Figure 19.3 shows the allocation differences between the unconditional Gaussian model and the conditional skewed t model for $A = 5$ (i.e. for $\varphi_v = 5/2$ and $\varphi_s = 5$).

The results show that the bulk of the difference between portfolios strategies depends largely on the use of different marginal distribution models. The first significant discrepancies arise when using time-varying Gaussian marginal distributions (CCC and DCC models) instead of unconditional Gaussian margins (Uncond. Gaussian model). In particular, the median position in equity resulting from the CCC model decreases significantly compared with the median position of the unconditional Gaussian strategy, especially for the lowest values of A . Accordingly, the median positions in commodity futures increases when we employ the CCC model. For example, for $A = 2$ ($\varphi_v = 1$ and $\varphi_s = 1$), the median positions of the CCC strategy are 0.50, 0.22 and -0.22 for oil, gold and equity, whereas for the unconditional model, the median positions are 0.36, 0.10 and 0.41, respectively. We also observe relevant differences between using Gaussian (CCC and DCC models) and generalized Student's t distributions (conditional copula models) for modelling the conditional margins. The main effect of introducing fat tails and asymmetry in the marginal distributions consists in an increase of long positions in gold and a reduction of long positions in oil. For instance, comparing the DCC and conditional Gaussian copula strategies for $A = 2$, we find that the median positions in oil and gold change from 0.43 and 0.21 to 0.21 and 0.46, respectively.

A second source of allocation differences is driven by the various types of dependence captured with our copula models. These discrepancies in optimal portfolio weights arise, first, from introducing a time-varying conditional dependence (e.g. CCC vs. DCC); and second, from considering tail dependence (e.g. t copula vs. Gaussian copula) and asymmetric dependence (e.g. skewed t vs. t copula). These allocation differences are significant mainly for $A = 1$ ($\varphi_v = 1/2$ and $\varphi_s = 1/3$) and $A = 2$ ($\varphi_v = 1$ and $\varphi_s = 1$). Specifically, when allowing for dynamic dependence, the median positions in the three assets decrease. For example, for $A = 1$, the CCC strategy generates median positions in oil, gold and equity equal to 0.88, 0.10 and -1.13 ; whereas for the DCC model the median weights are 0.74, 0.00 and -1.29 , respectively. When capturing asymmetric dependence using the skewed t copula, we also observe that median weights diminish with respect to the median positions corresponding to the symmetric t copula strategy. Thus, for $A = 1$, the median weights for the t copula model are 0.39, 0.57 and -0.82 ; and those obtained from the skewed t copula are equal to 0.28, 0.44 and -1.02 .

* Note that the short-sales constraint only affects the weights of spot contracts, in this case, the weights of the equity index.

TABLE 19.3 Summary Statistics of the Optimal Portfolio Weights for Different Strategies and Preferences

	Uncond. Gaussian			CCC			DCC			Gaussian Copula			<i>t</i> copula			skewed <i>t</i> copula		
	Oil	Gold	Equity	Oil	Gold	Equity	Oil	Gold	Equity	Oil	Gold	Equity	Oil	Gold	Equity	Oil	Gold	Equity
Panel A: Unconstrained																		
A = 1																		
5% pct.	0.00	0.00	-0.40	0.00	-0.54	-4.85	-0.12	-1.10	-5.26	0.00	-1.31	-5.65	0.00	-1.58	-5.43	-0.08	-1.57	-5.42
Median	0.60	0.00	0.18	0.88	0.10	-1.13	0.74	0.00	-1.29	0.29	0.49	-0.90	0.39	0.57	-0.82	0.28	0.44	-1.02
95% pct.	1.11	0.91	1.00	3.87	4.14	1.00	3.86	4.46	1.00	2.81	5.64	1.00	2.83	5.34	1.00	2.89	5.38	1.00
A = 2																		
5% pct.	0.09	0.00	0.03	-0.05	-0.40	-2.16	-0.18	-0.68	-2.42	-0.09	-0.75	-2.47	-0.03	-0.85	-2.45	-0.15	-0.89	-2.43
Median	0.36	0.10	0.41	0.50	0.22	-0.22	0.43	0.21	-0.30	0.21	0.46	-0.18	0.24	0.50	-0.13	0.20	0.45	-0.18
95% pct.	0.62	0.69	0.91	2.00	2.29	1.00	1.97	2.45	1.00	1.45	2.86	1.00	1.46	2.88	1.00	1.50	2.89	1.00
A = 5																		
5% pct.	0.10	0.09	0.29	-0.17	-0.30	-0.59	-0.23	-0.37	-0.65	-0.22	-0.40	-0.71	-0.20	-0.42	-0.67	-0.22	-0.43	-0.62
Median	0.19	0.33	0.45	0.26	0.33	0.27	0.24	0.31	0.27	0.14	0.41	0.27	0.15	0.40	0.26	0.14	0.42	0.26
95% pct.	0.31	0.56	0.76	0.88	1.21	1.00	0.86	1.29	1.00	0.67	1.39	0.97	0.65	1.39	1.00	0.67	1.42	1.00
A = 10																		
5% pct.	-0.05	0.29	0.38	-0.13	-0.24	-0.04	-0.21	-0.25	-0.06	-0.22	-0.29	-0.14	-0.20	-0.31	-0.12	-0.22	-0.29	-0.11
Median	0.12	0.40	0.46	0.19	0.31	0.40	0.17	0.34	0.41	0.11	0.41	0.37	0.12	0.38	0.38	0.11	0.42	0.38
95% pct.	0.19	0.56	0.60	0.49	0.87	0.80	0.50	0.91	0.80	0.42	1.00	0.84	0.39	0.98	0.84	0.40	0.98	0.86
Panel B: Short-sales Constrained																		
A = 1																		
5% pct.	0.00	0.00	0.00	0.00	0.00	0.00	-0.03	0.00	0.00	0.00	-0.33	0.00	0.00	-0.44	0.00	0.00	-0.34	0.00
Media	0.58	0.00	0.18	0.56	0.00	0.00	0.48	0.00	0.00	0.03	0.00	0.00	0.03	0.00	0.00	0.06	0.00	0.00
95% pct.	1.00	0.62	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
A = 2																		
5% pct.	0.09	0.00	0.03	-0.14	-0.13	0.00	-0.16	-0.15	0.00	-0.04	-0.42	0.00	-0.06	-0.34	0.00	-0.09	-0.50	0.00
Median	0.36	0.10	0.41	0.41	0.00	0.00	0.33	0.00	0.00	0.12	0.22	0.00	0.09	0.22	0.00	0.10	0.22	0.00
95% pct.	0.62	0.67	0.91	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
A = 5																		
5% pct.	0.10	0.09	0.29	-0.15	-0.30	0.00	-0.23	-0.31	0.00	-0.19	-0.38	0.00	-0.18	-0.42	0.00	-0.20	-0.43	0.00
Median	0.19	0.33	0.45	0.26	0.27	0.27	0.23	0.28	0.27	0.13	0.41	0.27	0.14	0.40	0.27	0.12	0.39	0.26
95% pct.	0.31	0.56	0.76	0.86	0.90	1.00	0.87	0.99	1.00	0.61	1.00	0.99	0.58	1.00	1.00	0.61	1.00	1.00
A = 10																		
5% pct.	-0.05	0.29	0.37	-0.11	-0.23	0.00	-0.21	-0.25	0.00	-0.21	-0.29	0.00	-0.19	-0.31	0.00	-0.21	-0.29	0.00
Median	0.12	0.40	0.46	0.19	0.32	0.41	0.16	0.33	0.40	0.11	0.40	0.38	0.12	0.38	0.39	0.11	0.42	0.38
95% pct.	0.19	0.54	0.60	0.49	0.80	0.81	0.50	0.85	0.81	0.42	0.91	0.83	0.39	0.91	0.84	0.40	0.93	0.87

Notes: This table reports the statistics for the allocation period related to the optimal portfolio weights for different multivariate models and preference specifications parameterized by the coefficient of relative risk aversion, A (i.e. $\varphi_V = A/2!$ and $\varphi_S = A(A+1)/3!$ in these cases). Panels A and B show the results for the unconstrained and short-sales constrained optimizations, respectively.




FIGURE 19.3 Optimal portfolio weights. Notes: Panel A shows the optimal portfolio weights over the out-of-sample period (June 2006 to September 2010) for the unconstrained conditional skewed t copula model for two preference specifications characterized by relative risk aversion coefficient $\mathcal{A} = 2$ and 5. Panel B shows the optimal portfolio weights for the unconstrained conditional skewed t copula model and the unconditional Gaussian model (a la Markowitz) for $\mathcal{A} = 5$.

Looking in more detail at the quantiles corresponding to our more flexible model, the conditional skewed t copula, we observe that the dispersion in the distribution of weights shrinks when increasing the value of $\mathcal{A}$ (see Figure 19.3). For the highest values of $\mathcal{A}$ considered, portfolio weights become less aggressive and there are less discrepancies among median positions of the skewed t strategy and those of less flexible models. Changes in the coefficient $\mathcal{A}$ affects the distribution of equity weights to a large extent. Thus, under the skewed t model, the median positions in equity for values of coefficient $\mathcal{A}$ equal to 1, 2, 5 and 10 are -1.02 , -0.18 , 0.26 and 0.38 , respectively; and the 5% percentiles are -5.42 , -2.43 , -0.62 and -0.11 . The median weights for oil futures range from 0.28 for $\mathcal{A} = 1$ to 0.11 for $\mathcal{A} = 10$; in this case, the decline in the median is mainly given by a decrease in the presence of extreme long positions in oil futures (the 95% percentile diminishes from 2.89 to 0.40). In contrast, increasing $\mathcal{A}$ generates less extreme long and short positions in gold, and the median weight varies slightly, from 0.44 for $\mathcal{A} = 1$ to 0.42 for $\mathcal{A} = 10$.

Panel B of Table 19.3 reports the percentiles of the distribution of portfolio weights when short sales of equity are not allowed. Imposing a short-sales constraint has a strong impact on the portfolio strategies, especially for the preference specifications corresponding to values of $\mathcal{A}$ equal to 1 and 2. Only the portfolio decisions obtained from the unconditional Gaussian model are not affected by this restriction. The models that yield median positions of being short in equity under the unconstrained portfolio optimization generate median positions equal to zero under the restricted allocation. As a result, the presence of extreme positions in commodity futures is reduced, especially the long positions in gold futures. For example, for $\mathcal{A} = 2$ the 5, 50 and 95% percentiles of the commodity positions resulting from the skewed t copula strategy under the short-sales constraint are -0.09 , 0.10 and 1.00 for oil, and -0.50 , 0.22 and 1.00 for gold; in contrast to the percentiles of the unrestricted portfolio decisions, which are -0.15 , 0.20 and 1.50 for oil, and -0.89 , 0.45 and 2.89 for gold. Therefore, for the least risk-averse investors under the short-sales restriction, part of the information included in the more flexible copula models is lost.*

* Patton (2004) reaches a similar conclusion for a portfolio selection problem with two stock indices.

19.4.3.2 Investment Ratios and Portfolio Performance

We now turn to analysing the moments, investment ratios and relative performance measures of the optimal portfolio returns resulting from the different copula strategies. Results are reported in Table 19.4 for various specifications of the three-moment preferences.

We compute the Sharpe, Sortino and Omega investment ratios, given, respectively, by

$$\text{Sharpe} = \frac{\mu_p - r_f}{\sigma_p}, \text{ Sortino} = \frac{\mu_p - r_f}{\sqrt{q_2^l(r_f)}}, \text{ and } \text{Omega} = \frac{q_1^u(r_f)}{q_1^l(r_f)}; \quad (19.18)$$

where μ_p is the average realized portfolio return, σ_p is the realized portfolio volatility, r_f is the risk-free rate, and $q_m^l(r_f)$ and $q_m^u(r_f)$ are the lower and upper partial moments of order m with target value equal to the risk-free rate.* The Sortino ratio modifies the Sharpe ratio by dividing the excess return of the portfolio by the downside standard deviation or square root of semi-variance. The Omega ratio can be interpreted as the probability weighted ratio of gains to losses, relative to the risk-free rate and it measures the combined effect of all of returns moments, rather than the individual effects of any of them. The higher the values of these ratios, the better the portfolio performance.

The Sharpe, Sortino and Omega ratios of the equally weighted portfolio are 0.05, 0.07 and 1.14 (see Table 19.4). The poor performance of equity markets and the boom of gold and oil during the 2006–2010 period reveals that these portfolios, with constant holdings in oil and gold futures, perform remarkably well compared with other strategies based on fitted distribution models. In particular, for almost all three-moment preferences considered, the investment ratios of the unconditional Gaussian model (Markowitz model) are substantially smaller than those corresponding to the equally weighted portfolio.

For the preferences considered, we find that, in general, the conditional copula models with generalized Student's t marginal distributions have larger investment ratios than the multivariate conditional Gaussian models (CCC and DCC). In addition, at least one of the t copula models (symmetric or skewed) perform better than the Gaussian copula, for most of the cases analysed. That is, conditional copulas that capture tail or asymmetric dependence, or both, achieve higher investment ratios across our allocation sample. For example, the Sortino and Omega ratios of the conditional skewed t copula strategy range from 0.10 and 1.19 to 0.14 and 1.28, whereas these investment ratios for the conditional Gaussian copula model range from 0.07 and 1.15 to 0.12 and 1.23. The DCC model yields Sortino and Omega ratios varying from 0.03 and 1.06 to 0.11 and 1.20; and the unconditional Gaussian strategy has ratios ranging from -0.004 and 0.99 to 0.07 and 1.14.

We also observe in Table 19.4 that portfolio decisions made using t and skewed t copulas have generally higher skewness coefficients than do Gaussian models. That is, taking into account tail dependence in the portfolio decisions could decrease the likelihood of negative portfolio returns.

Now, we compare the six model-driven portfolios in terms of their performance with respect to a benchmark strategy: the equally weighted portfolio. For that purpose, we employ two relative performance measures: the opportunity cost or performance fee (*Fee*) and the Graham–Harvey metric (GH). The performance fee is the amount that must be added to the return of the equally-weighted strategy, such that it leaves the investor indifferent to either portfolio decision. The Graham and Harvey (1997) measure is the difference between the alternative portfolio return and the volatility-matched benchmark portfolio. That is, to make both portfolios comparable in terms of volatility, we lever up/down the benchmark to match the alternative portfolio's volatility over the evaluation period. Table 19.4 reports both relative performance measures (in basis points per week) of the realized portfolio returns over the allocation period (2006–2010).

For all the preference specifications considered, both relative performance measures coincide in indicating the same best performing strategy (see values in bold in Table 19.4). These results suggest that

* The lower and upper partial moments of order m for a given target θ are defined as $q_m^l(\theta) = \int_{-\infty}^{\theta} (\theta - r)^m f_p(r) dr$ and $q_m^u(\theta) = \int_{\theta}^{\infty} (r - \theta)^m f_p(r) dr$, where $f_p(r)$ is the probability density function of the portfolio returns.

TABLE 19.4 Investment Ratios and Relative Performance Measures of the Realized Portfolio Returns.

	Unconstrained						Short-sales Constrained					
	Sharpe	Sortino	Omega	Skw	Fee	GH	Sharpe	Sortino	Omega	Skw	Fee	GH
Equally Weighted	0.050	0.071	1.142	−0.240	0.000	0.000	0.050	0.071	1.142	−0.240	0.000	0.000
Min. Variance	0.057	0.084	1.161	0.069	0.012	−0.047	0.057	0.084	1.161	0.069	0.012	−0.047
A = 1												
Uncond. Gaussian	0.004	0.005	1.011	−0.414	−0.189	−0.132	−0.003	−0.004	0.991	−0.275	−0.207	−0.160
GARCH-CCC	0.050	0.081	1.160	0.867	−0.005	0.364	0.042	0.064	1.116	0.186	−0.032	0.018
GARCH-DCC	0.044	0.071	1.138	0.857	−0.066	0.315	0.046	0.070	1.129	0.220	−0.016	0.036
Gaussian Copula	0.053	0.084	1.173	0.617	0.030	0.372	0.059	0.090	1.167	0.179	0.032	0.068
<i>t</i> Copula	0.061	0.100	1.205	1.191	0.102	0.457	0.045	0.068	1.124	0.152	−0.017	0.014
skewed <i>t</i> Copula	0.074	0.122	1.252	1.146	0.231	0.582	0.066	0.100	1.189	0.107	0.058	0.096
A = 2												
Uncond. Gaussian	0.018	0.024	1.051	−0.618	−0.098	−0.091	0.016	0.021	1.044	−0.629	−0.106	−0.099
GARCH-CCC	0.059	0.093	1.178	0.619	0.046	0.179	0.048	0.072	1.134	0.111	−0.007	0.030
GARCH-DCC	0.053	0.084	1.159	0.622	0.017	0.153	0.058	0.086	1.164	−0.018	0.029	0.065
Gaussian Copula	0.048	0.070	1.147	0.013	−0.015	0.110	0.067	0.100	1.188	−0.004	0.053	0.063
<i>t</i> Copula	0.083	0.154	1.296	2.375	0.196	0.353	0.063	0.095	1.174	0.145	0.040	0.050
skewed <i>t</i> Copula	0.078	0.127	1.249	0.861	0.147	0.265	0.069	0.102	1.191	−0.077	0.058	0.068
A = 5												
Uncond. Gaussian	0.042	0.058	1.119	−0.531	−0.021	−0.041	0.042	0.058	1.119	−0.531	−0.021	−0.041
GARCH-CCC	0.061	0.091	1.178	0.033	0.034	0.042	0.029	0.041	1.079	−0.212	−0.061	−0.065
GARCH-DCC	0.070	0.105	1.199	0.056	0.059	0.064	0.053	0.077	1.146	−0.165	0.008	−0.001
Gaussian Copula	0.078	0.117	1.228	0.046	0.078	0.073	0.059	0.086	1.167	−0.189	0.023	0.000
<i>t</i> Copula	0.072	0.108	1.208	0.109	0.061	0.056	0.066	0.097	1.188	−0.080	0.040	0.020
skewed <i>t</i> Copula	0.093	0.142	1.283	0.048	0.125	0.123	0.082	0.120	1.238	−0.262	0.080	0.060
A = 10												
Uncond. Gaussian	0.047	0.064	1.130	−0.506	−0.009	−0.036	0.051	0.071	1.144	−0.448	0.002	−0.025
GARCH-CCC	0.016	0.021	1.044	−0.493	−0.090	−0.107	0.018	0.025	1.050	−0.445	−0.083	−0.101
GARCH-DCC	0.040	0.057	1.111	−0.241	−0.024	−0.053	0.022	0.030	1.059	−0.292	−0.067	−0.096
Gaussian Copula	0.059	0.084	1.164	−0.191	0.019	−0.018	0.070	0.102	1.198	−0.226	0.043	0.005
<i>t</i> Copula	0.070	0.102	1.203	−0.244	0.045	0.012	0.058	0.083	1.160	−0.174	0.017	−0.018
skewed <i>t</i> Copula	0.091	0.135	1.269	−0.155	0.091	0.054	0.088	0.128	1.254	−0.268	0.080	0.041

Notes: This table reports the investment ratios and relative performance measures (in basis points per week) of the realized portfolio returns over the allocation (out-of-sample) period for different strategies. We first present the Sharpe, Sortino and Omega ratios, given respectively as $(\mu - r_f)/\sigma$, $(\mu - r_f)/\sqrt{q_2^l(r_f)}$, and $q_1^u(r_f)/q_1^l(r_f)$, where $q_m^u(\theta)$ and $q_m^l(\theta)$ are the upper and lower partial moments of order m for a given target θ . We also present the skewness (Skw) of the realized portfolio returns and two relative performance measures, the management fee (*Fee*) and the Graham–Harvey measure (*GH*), which are both calculated with respect to the equally weighted strategy. $\mathcal{A}$ is the coefficient of relative risk aversion for power utility functions. The highest value of each measure is marked in boldface.

the investor can obtain substantial gains using the portfolio rules based on conditional t copulas with generalized t marginal distributions for most of the preference specifications we report in Table 19.4. The gains seem to be higher for the skewed t copula portfolio for the unconstrained investor when $\mathcal{A}$ is lower, although we do not observe the same monotonic relation for the t or Gaussian copula portfolios or for the short-sales constrained investor. The opportunity costs of an investor holding the equally weighted portfolio instead of the portfolio based on the skewed t copula are between 3 and 12 basis points per year, whereas the opportunity costs for the Gaussian copula strategy range from -1 to 4 basis points per year. When we use the GH measure to compare alternative strategies with different levels of risk, we also find that copula models with tail dependence outperform Gaussian dependence models. In particular, the GH measure for the skewed t copula portfolios varies between 2 and 30 basis points per year; this measure for Gaussian copula portfolios ranges from 0 to 19 basis points; and for the DCC strategy, it varies between -5 and 16 basis points.

These results suggest that the allocation differences found among the multivariate copula portfolios imply also economic differences in terms of investment ratios and performance measures. The univariate higher moments and the tail and asymmetric dependence seem to be key features in this respect, especially for more aggressive investors. Specifically, as variance and loss aversion increases, skewed t copula strategies are less likely to produce large performance differences. These differences are also smaller for the short-sales constrained allocations, where the investor cannot increase their exposure to commodity futures by taking extreme short positions in equity. These results related to risk aversion and the performance of non-linear models is consistent with findings previously reported in the portfolio choice literature (e.g. Patton (2004) for the case of asymmetric dependence and wealth allocation between small and large cap stock indices; and Das and Uppal (2004) for the case of systemic risk and international portfolio choice). In the next section, we evaluate the robustness of our results, including in the analysis other parameterizations of the three-moment preferences.

19.4.3.3 Robustness Analysis

As a first robustness check, we investigate the allocation decisions made using three-moment preference specifications with arbitrary impacts of the portfolio variance and skewness (i.e. with arbitrary values of the coefficients φ_v and φ_s). Since the coefficients φ_v and φ_s may not necessarily be related to a third-order Taylor expansion of a given utility function, we have more flexibility to define ad-hoc three-moment investors, preferences; as in Harvey *et al.* (2010). In a separate supplemental data section, we report the percentiles of the distributions of optimal portfolio weights for some of these preference specifications. With respect to the allocation differences amongst our copula strategies, the results of this complementary approach are consistent with our previous findings. Furthermore, using arbitrary coefficients φ_v and φ_s we can do some sensitivity analysis and investigate the effect of increasing the impact of portfolio variance (φ_v) or skewness (φ_s) on the optimal allocation. More specifically, we find that when the impact of skewness (φ_s) increases, the median positions in equity and oil diminish, whereas long positions in gold increase.

For example, for $\varphi_v = 1/4$ and $\varphi_s = 0$, the skewed t copula model yields median positions equal to 0.34, 0.59 and -2.83 for oil, gold and equity; whereas if φ_s rises to $1/2$, the median positions are 0.09, 0.79 and -3.37. For $\varphi_v = 1/2$, increasing φ_s from $1/4$ to 1 modifies the median positions of oil, gold and equity from 0.28, 0.53 and -1.02 to 0.14, 0.54 and -1.18, respectively. On the other hand, when increasing the impact of variance (φ_v), the median and low quantiles of equity weights rise significantly, as we observed previously in Section 19.4.3.1. For instance, for $\varphi_s = 1/2$, when $\varphi_v = 1/4$, the skewed t strategy generates median positions equal to 0.09, 0.79 and -3.37; whereas if $\varphi_v = 1$, assets median positions are equal to 0.20, 0.45 and -0.17. In addition, the 5% quantile of equity weights increases from -10.0 to -2.39, and the 95% quantiles of oil and gold decrease from 5.33 and 11.0 to 1.50 and 2.84, respectively.

We also calculate the investment ratios and relative performance measures of the realized portfolio returns for these parameterizations of three-moment preferences. The conditional t copulas outperformed the multivariate Gaussian models in 4 of the 5 comparisons, suggesting again the importance

TABLE 19.5 Test for Superior Portfolio Performance.

	Unconstrained			Short-Sales Constrained			Harvey <i>et al.</i> (2010)		
	l	c	u	l	c	u	l	c	u
			A = 1			$\phi V = 1/4$ and $\phi S = 0$			
Equally weighted	0.037	0.049	0.049	0.232	0.232	0.251	0.104	0.104	0.104
Uncond. Gaussian	0.032	0.032	0.032	0.069	0.069	0.069	0.026	0.026	0.026
CCC	0.035	0.046	0.046	0.134	0.135	0.135	0.061	0.061	0.054
DCC	0.046	0.050	0.050	0.142	0.153	0.153	0.060	0.060	0.054
Gaussian Copula	0.137	0.192	0.192	0.499	0.713	0.713	0.706	0.858	0.547
<i>t</i> Copula	0.156	0.243	0.243	0.122	0.131	0.131	0.275	0.275	0.200
skewed <i>t</i> Copula	1.000	1.000	1.000	1.000	1.000	1.000	0.837	0.895	0.895
			A = 2			$\phi V = 1/4$ and $\phi S = 1/2$			
Equally weighted	0.086	0.113	0.113	0.141	0.152	0.152	0.170	0.170	0.162
Uncond. Gaussian	0.048	0.048	0.048	0.113	0.119	0.119	0.069	0.069	0.069
CCC	0.167	0.278	0.278	0.139	0.166	0.166	0.130	0.130	0.124
DCC	0.155	0.198	0.198	0.565	0.782	0.782	0.179	0.179	0.163
Gaussian Copula	0.109	0.111	0.111	0.888	0.894	0.894	0.167	0.167	0.089
<i>t</i> Copula	1.000	1.000	1.000	0.644	0.736	0.736	0.134	0.100	0.096
skewed <i>t</i> Copula	0.531	0.572	0.792	0.902	0.902	0.902	1.000	1.000	1.000
			A = 5			$\phi V = 1/2$ and $\phi S = 1/4$			
Equally weighted	0.128	0.135	0.135	0.131	0.133	0.133	0.137	0.137	0.128
Uncond. Gaussian	0.083	0.083	0.083	0.113	0.114	0.114	0.025	0.025	0.025
CCC	0.353	0.358	0.358	0.056	0.056	0.056	0.237	0.237	0.233
DCC	0.463	0.470	0.470	0.139	0.143	0.143	0.213	0.213	0.179
Gaussian Copula	0.418	0.564	0.564	0.155	0.155	0.155	0.270	0.270	0.215
<i>t</i> Copula	0.366	0.441	0.441	0.155	0.319	0.319	0.314	0.314	0.256
skewed <i>t</i> Copula	0.880	0.887	0.887	1.000	1.000	1.000	1.000	1.000	1.000
			A = 10			$\phi V = 1$ and $\phi S = 1/2$			
Equally weighted	0.142	0.153	0.157	0.135	0.143	0.148	0.261	0.261	0.249
Uncond. Gaussian	0.113	0.124	0.124	0.109	0.141	0.141	0.105	0.105	0.104
CCC	0.015	0.015	0.015	0.063	0.063	0.063	1.000	1.000	1.000
DCC	0.046	0.046	0.046	0.036	0.037	0.037	0.126	0.126	0.109
Gaussian Copula	0.111	0.111	0.111	0.131	0.214	0.235	0.107	0.107	0.107
<i>t</i> Copula	0.197	0.212	0.212	0.151	0.167	0.167	0.267	0.267	0.159
skewed <i>t</i> Copula	0.730	0.778	0.778	0.795	0.833	0.833	0.825	0.845	0.766

Notes: This table reports the reality check *p*-values of the Hansen (2005) test for superior portfolio performance for different benchmark models. Three *p*-values are reported: the consistent estimate (c), the upper (u) and lower (l) bounds. We employ the opportunity cost as the performance function in this case. For small *p*-values, we reject the hypothesis that the benchmark model performs as well as the best competing alternative model. Those *p*-values below 0.15 are marked in boldface. The implementation is based on the stationary bootstrap of Politis and Romano (1994).

of dependence specification for asset allocation with commodity futures. The GH measures (relative to the equally weighted portfolio) obtained from the skewed *t* copula model ranges between 14 and 86 basis points per year, while the DCC model vary between 8 and 46 basis points per year. See supplemental data for the complete set of results.

Finally, using the reality check test of Hansen (2005), we compare jointly the out-of-sample performance of the different portfolio strategies. To apply this superior predictive ability test, we first need to define a metric (loss function) and then employ this metric to compute the relative performance of each model with respect to a chosen benchmark model. The null hypothesis is that the benchmark is as good as any alternative model in terms of portfolio performance. That is, the test answers if any of

the models is better than the given benchmark. The stationary bootstrap of Politis and Romano (1994) is used to estimate the distribution of the test statistics under the null. Hansen (2005) proposed three p -values based on the test statistic estimates: a consistent p -value, as well as upper and lower bounds for the true p -value.

Employing the opportunity cost as our metric function (Patton 2004), we conduct the test of superior portfolio performance for different benchmark models. The corresponding p -values are reported in Table 19.5. The null hypothesis is rejected for small p -values. We observe that the probability of rejecting the null hypothesis for the equally weighted and unconditional Gaussian portfolios is higher (p -values < 0.10) for unconstrained strategies and for low values of $\mathcal{A}$, that is, for more aggressive investors. The DCC model is outperformed by other models, with p -values < 0.15 , in six of the twelve specifications we report. We are able to reject the conditional Gaussian copula model for three cases. By contrast, the conditional t copula strategy is only rejected twice, and there is little evidence that the conditional skewed t copula portfolio is outperformed by alternative models. Therefore, conditional copulas with tail dependence have generally superior out-of-sample performance for the different specifications considered.

19.5 Conclusion

This article investigates the portfolio selection problem of an investor with time-varying three-moment preferences when commodity futures are part of the investment opportunity set. In our specification, the portfolio returns' skewness provides a measure of the investor's loss aversion. We model the joint distribution of asset returns using a flexible multivariate copula setting that can disentangle the specific properties of each asset process from its dependence structure. The more general model we posit consists of a conditional skewed t copula with generalized Student's t marginal distributions and time-varying moments. Thus we can capture the specific distributional characteristics of commodity-futures returns and focus on their implications for the portfolio selection problem.

The empirical application employs weekly data for oil and gold futures and for the S&P 500 equity index, from June 1990 to September 2010. Our preliminary analysis and in-sample estimates suggest the presence of skewness and fat tails in the univariate processes, as well as evidence of both extreme and asymmetric dependence amongst oil, gold and equity. When computing the optimal portfolio weights, we find substantial discrepancies between the holdings obtained from our conditional copula models and those from more traditional Gaussian models. The key factors underlying these differences are the different specifications of the time-varying marginal distributions, the presence of dynamic conditional dependence among the univariate processes, and the modelling of tail and asymmetric dependence. The univariate higher moments and the type of tail dependencies are more relevant for aggressive investors. These discrepancies translate into economic differences in terms of better investment ratios and relative performance measures for the different specifications considered. For instance, depending on the allocation specification, the gains of using the conditional t copulas range up to 86 basis points per year for the period 2006–2010. The performance differences of portfolio strategies based on more flexible copulas are smaller when variance and loss aversion increase, as well as for short-sales constrained allocations. Furthermore, we analysis the robustness of our results, confirming that conditional copulas with tail dependence have generally superior out-of-sample performance for the different specifications considered.

Finally, some extensions to our analysis can be considered. For instance, it would be interesting to study the sensitivity of the investor's portfolio decisions to parameter uncertainty. Note that some cautions regarding the propagation of errors between the marginal distributions and the copula function have to be taken into account when implementing this type of analysis. Another possibility is to extend our portfolio selection problem with commodity futures to a dynamic asset allocation context. Thus, we could evaluate the hedging component of the optimal portfolio weights under the effects of skewness and asymmetric dependence. In addition, other possible extensions could include analysis of the impact of kurtosis as well as the impact of non-linear reward and risk measures, as in Biglova *et al.* (2010) and

the references therein. Furthermore, to gain more flexibility in the model specification, it would be of interest to consider alternative multivariate approaches with some non-parametric features. Some good candidates would be extensions of the semi-parametric copula-based models of Chen and Fan (2006) that would allow for time-varying copulas.

Acknowledgements

The authors are grateful to Michael Brennan, Álvaro Cartea, Claude Crampes, Jacinto Marabel, Rod McCrorie, Javier F. Navas, Eliezer Prisman, Michel Robe, Stefan Straetmans, and the seminar participants at INFINITI 2011, MFS 2011, FEBS 2012, IFABS 2012, EFMA 2012, Universidad Autónoma de Madrid, and Universidad Carlos III, for helpful comments and suggestions. M. Moreno acknowledges financial support by grants P08-SEJ-03917 and JCCM PPIII1-0290-0305.

Supplemental Data

Supplemental data for this article can be accessed at <http://dx.doi.org/10.1080/14697688.2014.935463>.

References

- Abramowitz, M. and Stegun, I.A., *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 1965 (Dover: New York).
- Ang, A. and Chen, J., Asymmetric correlations of equity returns. *J.Financ.Econ.*, 2002, **63**(3), 443–494.
- Barberis, N.C. and Huang, M., Stocks as lotteries: The implications of probability weighting for security prices. *Am.Econ.Rev.*, 2008, **98**(5), 2066–2100.
- Biglova, A., Ortobelli, S., Rachev, S.T. and Fabozzi, F., Modeling, estimation, and optimization of equity portfolios with heavy-tailed distributions. In *Optimizing Optimization: The Next Generation of Optimization Applications and Theory*, edited by S. Satchell, pp. 117–141, 2009 (Academic Press: London).
- Biglova, A., Ortobelli, S., Rachev, S.T. and Stoyanov, S., A note on the impact of non-linear reward and risk measures. *J.Appl.Funct.Anal.*, 2010, **5**(2), 194–202.
- Boerger, R.H., Cartea, Á., Kiesel, R. and Schindlmayr, G., Cross-commodity analysis and applications to risk management. *J.Futures Markets*, 2009, **29**(3), 197–217.
- Büyükaşahin, B., Haigh, M.S. and Robe, M.A., Commodities and equities: Ever a Market of One? *J.Altern.Invest.*, 2010, **12**(3), 76–95.
- Büyükaşahin, B. and Robe, M.A., Speculators, commodities, and cross-market linkages. *J.Int.Money Finance*, 2014, **42**, 38–70.
- Campbell, J.Y. and Hentschel, L., No news is good news: An asymmetric model of changing volatility in stock returns. *J.Financ.Econ.*, 1992, **31**(3), 281–318.
- Cartea, Á. and Figueroa, M., Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Appl.Math.Finance*, 2005, **12**(4), 313–335.
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J.Finance*, 2005, **60**(5), 2283–2331.
- Chen, X. and Fan, Y., Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification. *J.Econom.*, 2006, **135**(1), 125–154.
- Das, S.R. and Uppal, R., Systemic risk and international portfolio choice. *J.Finance*, 2004, **59**(6), 2809–2834.
- Daskalaki, C. and Skiadopoulos, G., Should investors include commodities in their portfolios after all? New evidence. *J.Bank.Finance*, 2011, **35**(10), 2606–2626.
- Demarta, S. and McNeil, A.J., The *t* copula and related copulas. *Int.Stat.Rev.*, 2005, **73**(1), 111–130.
- Diebold, F.X., Gunther, T.A. and Tay, A.S., Evaluating density forecasts with applications to financial risk management. *Int.Econ.Rev.*, 1998, **39**(4), 863–883.

- Dusak, K., Futures trading and investor returns: An investigation of commodity market risk premiums. *J.Polit.Econ.*, 1973, **81**(6), 1387–1406.
- Embrechts, P., Lindskog, F. and McNeil, A., Modelling dependence with copulas and applications to risk management. In *Handbook of Heavy Tailed Distributions in Finance*, edited by S.T. Rachev, pp. 329–384, 2003 (Elsevier/North-Holland: Amsterdam).
- Engle, R.F., Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J.Bus.Econ.Stat.*, 2002, **20**(3), 339–350.
- Engle, R.F. and Sheppard, K., Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. NBER Working Paper, 2001.
- Erb, C.B. and Harvey, C.R., The strategic and tactical value of commodity futures. *Financ.Anal.J.*, 2006, **62**(2), 69–97.
- Etula, E., Broker-dealer risk appetite and commodity returns. Federal Reserve Bank of New York Staff Reports, 406, 2010.
- Goffe, W.L., Ferrier, G.D. and Rogers, J., Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 1994, **60**(1–2), 65–99.
- Gorton, G. and Rouwenhorst, K.G., Facts and fantasies about commodity futures. *Financ.Anal.J.*, 2006, **62**(2), 47–68.
- Graham, J.R. and Harvey, C.R., Grading the performance of market-timing newsletters. *Financ.Anal.J.*, 1997, **53**(6), 54–66.
- Guidolin, M. and Timmermann, A., International asset allocation under regime switching, skew, and kurtosis preferences. *Rev.Financ.Stud.*, 2008, **21**(2), 889–935.
- Hansen, B.E., Autoregressive conditional density estimation. *Int.Econ.Rev.*, 1994, **35**(3), 705–730.
- Hansen, P.R., A test for superior predictive ability. *J.Bus.Econ.Stat.*, 2005, **23**(4), 365–380.
- Härdle, W.K., Okhrin, O. and Okhrin, Y., Time varying hierarchical archimedean copulae. Discussion Paper SFB 649, Humboldt University, 2010.
- Harvey, C.R., Liechty, J.C., Liechty, M.W. and Müller, P., Portfolio selection with higher moments. *Quant.Finance*, 2010, **10**(5), 469–485.
- Harvey, C.R. and Siddique, A., Autoregressive conditional skewness. *J.Financ.Quant.Anal.*, 1999, **34**(4), 465–487.
- Harvey, C.R. and Siddique, A., Conditional skewness in asset pricing tests. *J.Finance*, 2000, **55**(3), 1263–1295.
- Hong, H. and Yogo, M., What does futures market interest tell us about the macro-economy and asset prices? *J.Financ.Econ.*, 2012, **105**(3), 473–490.
- Irwin, S.H. and Sanders, D.R., Index funds, financialization, and commodity futures markets. *Appl.Econ.Perspect.Policy*, 2011, **33**(1), 1–31.
- Jondeau, E. and Rockinger, M., Conditional volatility, skewness, and kurtosis: Existence, persistence, and comovements. *J.Econ.Dyn.Control*, 2003, **27**(10), 1699–1737.
- Jondeau, E. and Rockinger, M., The copula-garch model of conditional dependencies: An international stock market application. *J.Int.Money Finance*, 2006, **25**(5), 827–853.
- Jondeau, E. and Rockinger, M., On the importance of time variability in higher moments for asset allocation. *J.Financ.Econom.*, 2012, **10**(1), 84–123.
- Kat, H.M. and Oomen, R.C.A., What every investor should know about commodities. Part II: multivariate return analysis. *J.Invest Manage.*, 2007, **5**(3), 16–40.
- Kraus, A. and Litzenberger, R.H., Skewness preference and the valuation of risk assets. *J.Finance*, 1976, **31**(4), 1085–1100.
- Longin, F. and Solnik, B., Extreme correlation of international equity markets. *J.Finance*, 2001, **56**(2), 649–676.
- Manner, H. and Reznikova, O., A survey on time-varying copulas: Specification, simulations, and application. *Econom.Rev.*, 2012, **31**(6), 654–687.
- Mardia, K.V., Measures of multivariate skewness and kurtosis with applications. *Biometrika*, 1970, **57**(3), 519–530.

- McNeil, A.J., Frey, R. and Embrechts, P., *Quantitative Risk Management: Concepts, Techniques, and Tools*, 2005 (Princeton University Press: Princeton, NJ).
- Mitton, T. and Vorkink, K., Equilibrium underdiversification and the preference for skewness. *Rev. Financ. Stud.*, 2007, **20**(4), 1255–1288.
- Nelsen, R.B., *An Introduction to Copulas*, 2006 (Springer: New York).
- Ohana, S., Modeling global and local dependence in a pair of commodity forward curves with an application to the us natural gas and heating oil markets. *Energy Econ.*, 2010, **32**(2), 373–388.
- Ortobelli, S., Biglova, A., Rachev, S.T. and Stoyanov, S., Portfolio selection based on a simulated copula. *J. Appl. Funct. Anal.*, 2010, **5**(2), 177–193.
- Patton, A.J., On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *J. Financ. Econom.*, 2004, **2**(1), 130–168.
- Patton, A.J., Estimation of multivariate models for time series of possibly different lengths. *J. Appl. Econom.*, 2006a, **21**(2), 147–173.
- Patton, A.J., Modelling asymmetric exchange rate dependence. *Int. Econ. Rev.*, 2006b, **47**(2), 527–556.
- Politis, D.N. and Romano, J.P., The stationary bootstrap. *J. Am. Stat. Assoc.*, 1994, **89**(428), 1303–1313.
- Routledge, B.R., Seppi, D.J. and Spatt, C.S., Equilibrium forward curves for commodities. *J. Finance*, 2000, **55**(3), 1297–1338.
- Samuelson, P.A., The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. *Rev. Econ. Stud.*, 1970, **37**(4), 537–542.
- Sklar, A., Fonctions de répartition à n dimensions et leurs marges [Distribution functions with n dimensions and their margins]. *Publications de l'Institut Statistique de l'Université de Paris* [Publ. Inst. Statistique Univ. Paris], 1959, **8**, 229–231.
- Tang, K. and Xiong, W., Index investment and financialization of commodities. *Financ. Anal. J.*, 2012, **68**(6), 54–74.
- Zolotko, M. and Okhrin, O., Modelling general dependence between commodity forward curves. Discussion Paper SFB 649, Humboldt University, 2012.

Appendix A: Copula Functions

This section describes the three implicit copulas we propose as dependence functions for our multivariate model: the Gaussian, t , and skewed t copulas (see [Section 19.3](#)). These three copulas correspond to the dependence functions contained in three multivariate normal mixture distributions (see McNeil *et al.* (2005)). This class of distributions adopts the following representation:

$$\mathbf{X} = \boldsymbol{\mu} + W\boldsymbol{\gamma} + \sqrt{W}\mathbf{Z} \quad (\text{A.1})$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ are parameter vectors in $\mathbb{R}^d$, $\mathbf{Z} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, and W is a random variable independent of $\mathbf{Z}$. When the mixing random variable W satisfies $\eta/W \sim \chi^2(\eta)$, then the resulting mixture distribution of the random vector $\mathbf{X}$ is termed the asymmetric or skewed t distribution $H(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\eta}, \boldsymbol{\gamma})$, which belongs to the wider family of multivariate generalized hyperbolic distributions.

Applying Sklar's theorem in Equation (19.6), we can obtain the skewed t copula function from the generalized hyperbolic skewed t distribution $H(\mathbf{0}, \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma})$ defined by $\boldsymbol{\mu} = \mathbf{0}$ and the correlation matrix $\mathbf{P}$ implied by the dispersion matrix $\boldsymbol{\Sigma}$.^{*} Then the skewed t copula is defined as

$$C^{\text{SK}}(\mathbf{u}; \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma}) = H(H_1^{-1}(u_1; \boldsymbol{\eta}, \boldsymbol{\gamma}_1), \dots, H_d^{-1}(u_d; \boldsymbol{\eta}, \boldsymbol{\gamma}_d); \mathbf{P}, \boldsymbol{\eta}, \boldsymbol{\gamma}), \quad (\text{A.2})$$

where the $H_i(\cdot; \boldsymbol{\eta}, \boldsymbol{\gamma}_i)$ are the d univariate skewed t distribution functions, the H_i^{-1} are the corresponding quantile functions, and $\mathbf{u} = (u_1, \dots, u_d)'$ is the probability transformed vector.

^{*} The copula function is invariant under any strictly increasing transformation of the marginal distributions, including the standardization of the components of the random vector $\mathbf{X}$.

Special cases can be obtained from the normal mixture representation in Equation (A.1). When $\gamma = \mathbf{0}$, we have the multivariate Student's t distribution; obviously, when W is constant we obtain the multivariate Gaussian distribution. Thus, the unique t copula of a d -variate Student's t distribution can be expressed as

$$C^T(\mathbf{u}; \mathbf{P}, \eta) = T(T^{-1}(u_1; \eta), \dots, T^{-1}(u_d; \eta); \mathbf{P}, \eta), \quad (\text{A.3})$$

where $T(\cdot; \mathbf{P}, \eta)$ is the joint distribution function of a d -variate Student's t distribution with η degrees of freedom and correlation matrix $\mathbf{P}$ and $T^{-1}(u_i; \eta)$ is the inverse function of the univariate Student's t distribution with η degrees of freedom. In the same way, we can define the d -variate Gaussian copula as

$$C^G(u; \mathbf{P}) = \Phi(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d); \mathbf{P}), \quad (\text{A.4})$$

where $\Phi(\cdot; \mathbf{P})$ denotes the joint distribution function of the d -variate standard normal distribution with correlation matrix $\mathbf{P}$ and Φ^{-1} denotes the inverse of the univariate standard normal distribution.

We proceed to compute the density functions of the three copulas. The density function of a parametric copula that is absolutely continuous is given by

$$c(\mathbf{u}) = \frac{\partial^d C(u_1, \dots, u_d)}{\partial u_1 \cdots \partial u_d}. \quad (\text{A.5})$$

For the case of the three implicit copulas we consider here, the density functions can be obtained from differentiating Equations (A.2)–(A.4). Thus, the density function of the d -variate skewed t copula can be expressed as

$$c^{\text{SK}}(\mathbf{u}; \mathbf{P}, \eta, \gamma) = \frac{h(H_1^{-1}(u_1; \eta, \gamma_1), \dots, H_d^{-1}(u_d; \eta, \gamma_d); \mathbf{P}, \eta, \gamma)}{h(H_1^{-1}(u_1; \eta, \gamma_1); \eta, \gamma_1) \cdots h_d(H_d^{-1}(u_d; \eta, \gamma_d); \eta, \gamma_d)} \quad (\text{A.6})$$

where $h(\cdot; \mathbf{P}, \eta, \gamma)$ is the joint density of the multivariate skewed t distribution H , and the $h_i(\cdot; \eta, \gamma_i)$ are its corresponding marginal density functions. Using the results from McNeil *et al.* (2005, Section 3.2.3) for the density functions of generalized hyperbolic distributions, and some algebra, we explicitly obtain the density function of the d -variate skewed t copula, given by

$$c^{\text{SK}}(\mathbf{u}; \mathbf{P}, \eta, \gamma) = \frac{K_{\frac{\eta+d}{2}} \left(\sqrt{(\eta + \mathbf{x}' \mathbf{P}^{-1} \mathbf{x}) \gamma' \mathbf{P}^{-1} \gamma} \right) \left(\sqrt{(\eta + \mathbf{x}' \mathbf{P}^{-1} \mathbf{x}) \gamma' \mathbf{P}^{-1} \gamma} \right)^{\frac{\eta+1}{2}} e^{\mathbf{x}' \mathbf{P}^{-1} \gamma}}{\prod_{i=1}^d K_{\frac{\eta+1}{2}} \left(\sqrt{(\eta + x_i^2) \gamma_i^2} \right) \left(\sqrt{(\eta + x_i^2) \gamma_i^2} \right)^{\frac{\eta+d}{2}} e^{x_i \gamma_i}} \quad (\text{A.7})$$

$$\times |\mathbf{P}|^{-1/2} \left(\frac{\Gamma(\frac{\eta}{2})}{21 - \eta/2} \right)^{d-1} \frac{\prod_{i=1}^d \left(1 + \frac{x_i^2}{\eta} \right)^{(\eta+1)/2}}{\left(1 + \frac{\mathbf{x}' \mathbf{P}^{-1} \mathbf{x}}{\eta} \right)^{(\eta+d)/2}},$$

where $\mathbf{x} = (x_1, \dots, x_d)'$ and $x_i = H_i^{-1}(u_i; \eta, \gamma_i)$. In addition, $\Gamma(\cdot)$ denotes the Gamma function and K_ρ is the modified Bessel function of the second kind with order ρ , which can be implemented numerically (see Abramowitz and Stegun (1965) for more details about these functions and their properties). A similar expression to that in Equation (A.5) can be derived for the t and Gaussian copulas from their respective joint and marginal density functions. Thus the density function of the d -variate t copula is given by

$$c^T(\mathbf{u}; \mathbf{P}, \eta) = |\mathbf{P}|^{-1/2} \frac{\Gamma\left(\frac{\eta+d}{2}\right) \Gamma\left(\frac{\eta}{2}\right)^{d-1} \prod_{i=1}^d \left(1 + \frac{x_i^2}{\eta}\right)^{(\eta+1)/2}}{\Gamma\left(\frac{\eta+1}{2}\right)^d \left(1 + \frac{\mathbf{x}' \mathbf{P}^{-1} \mathbf{x}}{\eta}\right)^{(\eta+d)/2}}, \quad (\text{A.8})$$

where $x_i = T^{-1}(u_i; \eta)$. Finally, the density function of the d -variate Gaussian copula is expressed as

$$c^G(\mathbf{u}; \mathbf{P}) = |\mathbf{P}|^{-1/2} \exp\left(-\frac{1}{2} \mathbf{x}'(\mathbf{P}^{-1} - \mathbb{I}_d) \mathbf{x}\right), \quad (\text{A.9})$$

where $x_i = \Phi^{-1}(u_i)$ and $\mathbb{I}_d$ denotes the unit matrix of size d .

Appendix B: The Two-Stage Log-Likelihood Function

Following Nelsen (2006, Theorem 2.10.9) and Patton (2006b, Theorem 1), Equation (19.6) presents a multivariate and conditional extension to Sklar's theorem. Then, according to Equation (19.6), the conditional density function of the joint distribution $F_t(r_{1,t+1}, \dots, r_{d,t+1}; \boldsymbol{\theta})$ is given by

$$f_t(r_{t+1}; \boldsymbol{\theta}) = \frac{\partial^d F_t(r_{t+1}; \boldsymbol{\theta})}{\partial r_{1,t+1} \cdots \partial r_{d,t+1}} = \prod_{i=1}^d f_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_i, M) \cdot c_t(u_{1,t+1}, \dots, u_{d,t+1}; \boldsymbol{\theta}_C), \quad (\text{B.1})$$

where the $u_{i,t+1} = F_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_{i,M})$ are the marginal conditional distributions; $f_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_{i,M})$ are the marginal conditional density functions; and $c_t(u_{1,t+1}, \dots, u_{d,t+1}; \boldsymbol{\theta}_C)$ is the conditional copula density function (defined in Equation (A.5)).

Taking logarithms in Equation (B.1) and summing for all observations in the sample, $\bar{\mathbf{r}}_T = \{\mathbf{r}_1, \dots, \mathbf{r}_T\}$, we determine that the log-likelihood function of the joint model $\mathcal{L}(\boldsymbol{\theta}; \bar{\mathbf{r}}_T)$ in Equation (19.16) can be divided in two terms, as follows:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}; \bar{\mathbf{r}}_T) &= \sum_{i=1}^d \sum_{t=1}^T \log f_{i,t}(r_{i,t+1}; \boldsymbol{\theta}_i, M) + \sum_{t=1}^T \log c_t(u_{1,t+1}, \dots, u_{d,t+1}; \boldsymbol{\theta}_C) \\ &= \sum_{i=1}^d \mathcal{L}_i(\boldsymbol{\theta}_{i,M}; \bar{\mathbf{r}}_T) + \mathcal{L}_C(\boldsymbol{\theta}_C; \boldsymbol{\theta}_M, \bar{\mathbf{r}}_T), \end{aligned} \quad (\text{B.2})$$

where $\mathcal{L}_i$ and $\mathcal{L}_C$ are the log-likelihood functions for the i th marginal model and for the copula function, respectively. Moreover, $\boldsymbol{\theta}_M = (\boldsymbol{\theta}_{1,M}, \dots, \boldsymbol{\theta}_{d,M})'$ is the parameter set of the d marginal conditional distributions, and $\boldsymbol{\theta}_C$ is the parameter set of the conditional copula.

Strategic Commodity Allocation

20.1	Introduction	433
20.2	The Optimal Strategies Design.....	435
20.3	Empirical Assessment of the Strategies.....	440
	Data Description • Parameters and State Variables Inference • Optimal Strategies	
20.4	Concluding Remarks.....	448
	References.....	452
	Appendix A: The Volatilities of the Zero-Coupon Bond and Futures Price	454
	Appendix B: The Computation of the Certainty Equivalent.....	454
	Appendix C: (Log) Likelihood Functions	456

Pierre Six

This article extends the study of the financialization of commodities (Rouwenhorst and Tang [Annu. Rev. Financ. Econ., 2012, 4, 447–467]) by considering an investment in the term structure of commodity futures prices. Specifically, we analyse the benefits of adding a distant commodity futures contract and/or a spot commodity (near month futures contract) to a portfolio of bonds and stocks in a setting similar to Brennan and Schwartz [The use of treasury bill futures in strategic asset allocation programs. In *Worldwide Asset and Liability Modeling*, edited by W.T. Ziemba and J.M. Mulvey, pp. 205–230, 1998 (Cambridge University Press: Cambridge)]. Our analysis employs an empirical study that covers the post-financial crisis period. We show that the spot commodity considerably improves the value of the portfolio. However, an investment in the whole term structure of futures contracts is optimally achieved through high opposite positions in the spot commodity and distant futures contracts. We find that these extreme calendar spreads can result in an inappropriate investment.

Keywords: Commodity investment; Commodity financialization; Dynamic asset allocation; Commodity risk premium; Theory of storage

20.1 Introduction

The large inflow of money invested in commodities in the past decades as described by Cheng and Xiong (2013) has recently led researchers to focus on the financialization process of commodity markets (Rouwenhorst and Tang 2012, Singelton 2012, Tang and Xiong 2012, Basak and Pavlova 2013). These papers analyse empirically and/or theoretically the impact of this financialization process on the change in the properties of commodities prices. However, these manuscripts do not focus on commodity investment. This paper studies an investment in the term structure of commodity prices with a data-set of commodity futures prices subsequent to this financialization process.

The recent allocation studies (e.g. Erb and Harvey 2006; Miffre and Rallis 2007, Marshall *et al.* 2008, Fuertes *et al.* 2010, Basu and Miffre 2013) envisage a tactical allocation because of the strong consensus regarding the significance of a commodity conditional risk premium.* However, these studies tend to focus on ex post statistical frameworks that rely on the hedging pressure theory† or on ad hoc trading rules. We focus on strategic asset allocation as defined by Brennan and Schwartz (1998) in the case of bonds and stocks. We rely on the theory of storage (Working 1948, 1949) to model the information conveyed by the commodity futures market.‡

Brennan and Schwartz (1998) show the advantage of a dynamic investment strategy employing futures contracts that uses the predictability of asset returns for the equity and bond markets, as opposed to using futures to hedge underlying positions. The time varying nature of risk premia in commodity markets is justified by the predictability of commodity prices (Bessembinder and Chan 1992, Gorton *et al.* 2007, Hong and Yogo 2010). Our study then extends the analysis of Brennan *et al.* (1997) to commodity markets. To do so, we rely on the seminal works of Black (1976) and Gibson and Schwartz (1990), who adapt the theory of storage to the continuous time arbitrage free evaluation framework. We specifically use an extension similar to that of Casassus and Collin-Dufresne (2005), which highlights the relevance of time-varying commodity risk premia for the valuation of commodity derivatives. Their framework uses the short rate, the spot commodity price and the convenience yield as predictive factors/state variables for the term structure of the commodity futures prices as well as for the term structure of interest rates. For equity, we use a simplified version of the model of Wachter (2002) with an affine market price of risk of the (log) equity price. The time-varying risk premium in the bond market is also considered in our setting as an affine function of the short rate (Dai and Singleton 2002, Duffee 2002).

The dynamics of the returns on assets in our setting are affine functions of (conditionally) Gaussian predictive factors. We study a pure allocation problem where our investor maximizes utility stemming only from terminal wealth. His/her relative risk aversion is constant as advocated by empirical studies (Meyer and Meyer 2005). As a consequence, the computation of the optimal strategies necessitates only the computation of ordinary differential equations even when an investment in the short end of the term structure is considered (Liu 2007).

The bond-stock mix allocation provides a natural benchmark for our study: each investor trades at least bonds and equity. We consider two embedded commodity investment scenarios. In the first one, the investor takes full advantage of the information provided by the commodity term structure of prices: he/she invests in both the spot commodity§ and in a (distant) futures contract written on this commodity. The second scenario is represented by an investor who focuses on the information provided by the spot price and invests only in the spot commodity, i.e. the short end of the term structure.

* Bessembinder (1992), Bessembinder and Chan (1992), Miffre (2000), DeRoos *et al.* (2000), Scherer and He (2008) and Willenbrock (2011) are examples of articles that discuss the existence of a conditional commodity risk premium.

† The reader can refer to Bessembinder (1992) for a modern account of the hedging pressure theory.

‡ Related studies include Hong (2001), Dai (2009) and Su and Keung (2010). Hong (2001) considers an investment both in the spot commodity and in a distant futures contract under the theory of storage. However, he does not consider an investment in bonds and equity. Moreover, he focuses on the impact of market clearing on the open interest of the commodity futures markets. Dai (2009) and Su and Keung (2010) do consider an investment in bonds and equity. However, Dai (2009) considers a constant interest rate and a constant equity opportunity set, while Su and Keung (2010) use a vector autoregressive approach, which results in a modelling of bonds, equity and commodities that is purely statistical. As a consequence, the theory of storage or other economic theories related to commodities are not apparent in their frameworks. Additionally, Dai (2009) and Su and Keung (2010) limit their study to multi-commodities indexes. Investments in single commodities are nevertheless considerable as shown by the monthly reports of the single commodity indexes released by the Commodity Futures Trading Commission (CFTC).

§ While a direct and frictionless investment in the spot commodity might not be available, our setting should be considered to be an idealized case of a short-term futures contract investment. Some articles nevertheless argue in favour of the possibility of a direct investment in the physical commodity (e.g. Fung and Hsieh 1997).

This investor then neglects the information given by the basis.* For each scenario, we compare two embedded cases: a constant market price of risk and a dynamic market price of risk as described above. We implement our theoretical results using market data that includes three weekly economically related commodities: copper, gold and oil.

We focus on two different ex post analyses: an analysis of the Sharpe ratio as in Basu and Miffre (2013) and an analysis of the terminal wealth, as in Brennan *et al.* (1997). Both analyses rule out the benefit of investing in a distant futures contract in addition to the spot commodity. This optimal investment leads to very high positions in the spot commodity and the distant futures contract that are opposite in sign. These extreme calendar spreads combined with the high correlation between the spot commodity and the distant futures contract prices can result in inappropriate investments. While the Sharpe ratio analysis provides mixed evidence, the analysis of the terminal wealth clearly demonstrates the benefit of investing in the spot commodity in addition to bonds and equity. Our study also shows the benefit of considering a dynamic market price of risk. Our results are consistent across investors' risk aversions and investment horizons.

The remainder of the paper is organized as follows. Section 20.2 describes the economic framework and derives the optimal strategies. Section 20.3 is an empirical study of these commodity strategies for the copper, gold and oil markets. Section 20.4 offers concluding remarks. The mathematical proofs are available from the author upon request.

20.2 The Optimal Strategies Design

Uncertainty in our economy is described by a complete filtered probability space $(\Omega; \mathcal{F}; \Theta; P)$ that is endowed with a continuous non-decreasing filtration $\Theta \equiv \{F_t; t \in [0; T]\}$. T is a positive constant that represents the end of the economy and $F_T \equiv F$. The processes $z_{Sp}, z_{rP}, z_{XP}, z_{\delta P}, t \in [0; T]$ are correlated standard Brownian motions defined on $(\Omega, \mathcal{F}, \Theta, P)$ that represent the risks of the equity price, the short rate, the commodity spot price and the commodity convenience yield, respectively. $\Theta := \{F_t; t \in [0; T]\}$ can then be understood as the augmented filtration generated by the paths of these Brownian motions. The instantaneous correlation between two Brownian motions is constant: $\rho_{ij} dt \equiv_{\text{CE}} dz_{it} dz_{jt}, i, j \in \{S, r, X, \delta\}$, where $\{S, r, X, \delta\}$ designate the equity price, the short interest rate, the commodity spot price and its convenience yield, respectively. For the remainder of the article, “ $\cdot$ ” stands for the transpose symbol and E_t represents the expectation operator conditional on F_t .

We rely on the models of Casassus and Collin-Dufresne (2005) and of Wachter (2002) to lay out the dynamics of our market. However, given the number of factors that we consider in our study, we rely on a simplified version of these models to keep the analysis tractable. Our simultaneous exogenous definition of the commodity spot prices, convenience yields, interest rates and equity dynamics is motivated by the fact that our investor only trades one commodity in addition to bonds and equity in our partial equilibrium. The dynamics under consideration are summarized by the following system of stochastic differential Equations (20.1a)–(20.1d):

$$\frac{dS_t}{S_t} = [r_t + \beta_{s0} + \beta_{ss}s_t]dt + \sigma_s dz_{St}, s_t = \log(S_t), \quad (20.1a)$$

$$dr_t = \kappa_r [\theta_r - r_t]dt + \sigma_r dz_{rt}, \quad (20.1b)$$

$$\begin{aligned} \frac{dX_t}{X_t} + \delta_t dt &= [r_t + \beta_{X0} + \beta_{Xs}x_t + \beta_{X\delta}\delta_t]dt + \sigma_X dz_{Xt}, x_t \\ &\equiv \ln(X_t), \end{aligned} \quad (20.1c)$$

$$d\delta_t = \kappa_\delta [\theta_\delta - \delta_t]dt + \sigma_\delta dz_{\delta t}. \quad (20.1d)$$

* The basis is usually defined as the difference between a distant futures price and the spot price.

The dynamics of the equity price (Equation (20.1a)) are described by a constant volatility σ_s and an affine market price of risk, $\lambda_{st} = [\beta_{s0} + \beta_{ss}s_t]/\sigma_s$. The short interest rate (Equation (20.1b)) follows a mean-reverting process with constant speed of mean-reversion, long-term mean and volatility, κ_r , θ_r , σ_r , respectively (Vasicek 1977). The short interest rate market price of risk is an affine function of itself: $\lambda_{rt} = [\beta_{r0} + \beta_{rr}r_t]/\sigma_r$ (Duffee 2002, Dai and Singleton 2002, Casassus and Collin-Dufresne 2005). The commodity spot price, X_t , exhibits a constant volatility, σ_x , and an affine market price of risk equal to $\lambda_{xt} \equiv [\beta_{x0} + \beta_{xx}x_t + \beta_{x\delta\delta t}]/\sigma_x$. Its convenience yield, δ_t , has a constant volatility, σ_δ , a constant speed of mean-reversion, κ_δ , and a long-term mean θ_δ (Schwartz 1997). Finally, the market price of risk of the convenience yield is an affine function of itself: $\lambda_{\delta t} = [\beta_{\delta 0} + \beta_{\delta\delta}\delta_t]/\sigma_\delta$. The risk neutral speed of mean-reversion and long-term mean of the convenience yield are then given as follows: $\kappa_{\delta Q} = \kappa_\delta + \beta_{\delta\delta}$, $\theta_{\delta Q} = [\kappa_\delta\theta_\delta - \beta_{\delta 0}]/\kappa_{\delta Q}$. The constant market prices of risk are obtained by setting $\beta_{ss} = \beta_{rr} = \beta_{xx} = \beta_{x\delta} = \beta_{\delta\delta} = 0$.

We denote by $P_t(T_p)$ the time- t price of a (unit) zero-coupon bond maturing at T_p and by $F_t(T_F)$ the time- t futures price of the delivery of a unit of the spot commodity at maturity T_F . Using the standard asset pricing arguments (available from the author upon request), the dynamics of these two risky assets are given as follows:

$$\frac{dP_t(T_p)}{P_t(T_p)} = [r_t - \sigma_p(t; T_p)\lambda_{rt}]dt - \sigma_p(t; T_p)dz_{rt}, \quad (20.2a)$$

$$\begin{aligned} \frac{dF_t(T_F)}{F_t(T_F)} = & [\sigma_{Fr}(t; T_F)\lambda_{rt} + \sigma_{Fx}(t; T_F)\lambda_{xt} + \sigma_{F\delta}(t; T_F)\lambda_{\delta t}]dt \\ & + \sigma_{Fr}(t; T_F)dz_{rt} + \sigma_{Fx}(t; T_F)dz_{xt} + \sigma_{F\delta}(t; T_F)dz_{\delta t}, \end{aligned} \quad (20.2b)$$

where the deterministic functions of time, $\sigma_p(t; T_p)$, $\sigma_{Fr}(t; T_F)$, $\sigma_{Fx}(t; T_F)$, and $\sigma_{F\delta}(t; T_F)$, are provided in Appendix A.

For all of our embedded allocation issues, the investor can invest in the risk-free asset $N_t \equiv \exp\left(\int_0^t r_u du\right)$. Moreover, we know from Merton (1969, 1971, 1973) that in our frictionless market, the investment in the risk-free market can be deduced from that in the risky assets. As a consequence, we focus only on the allocation in the risky markets. First, we consider an investor who only trades in the bond and equity markets. Every quantity specific to this benchmark investment will be denoted with a superscript '(a)'. Second, we consider an investor who can trade bonds and equity as well as the spot commodity, X_t , and a (distant) futures contract, $F_t(T_F)$. Each quantity related to this strategy will be denoted with a superscript '(b)'. The case where the investor trades only the spot commodity X_t in addition to equity and bonds is denoted by '(β)'.

The investor's constant relative risk aversion towards wealth is denoted by γ and his/her investment horizon by T_I . Following Merton (1969, 1971, 1973), we can show that the investor faces the following investment programs, $P_{ro}^{(n)}, n \in \{a, b, \beta\}$:

$$P_{ro}^{(n)} : J_t^{(n)} \equiv \sup_{\pi_t^{(n)}, t \leq u \leq T_I} E_t \left[u(W_{T_I}^{(n)}) \right], u(x) \equiv \frac{x^{1-\gamma}}{1-\gamma}, \quad (20.3a)$$

$$\frac{dW_t^{(n)}}{W_t^{(n)}} = \left[r_t + \left(\sigma_A^{(n)}(t) \pi_t^{(n)} \right)' \lambda_t^{(n)} \right] dt + \left(\sigma_A^{(n)}(t) \pi_t^{(n)} \right)' dz_t^{(n)}, \quad (20.3b)$$

where $W_t^{(n)}$ designates the time- t wealth of the investor who invests in program $P_{ro}^{(n)}$. $\sigma_A^{(n)}(t)$, $\pi_t^{(n)}$, $\lambda_t^{(n)}$, $z_t^{(n)}$ designate the volatility matrix of the risky assets, the vector of the wealth proportions invested in the risky assets, the vector of the market prices of risks of traded assets and the vector of the Brownian motions necessary to compute the program $P_{ro}^{(n)}$, respectively. For later reference, we denote by

$\Sigma_A^{(n)}(t) \equiv \sigma_A^{(n)}(t)' \sigma_A^{(n)}(t)$ the variance-covariance matrix of the traded assets, and by $\mu_{At}^{(n)}$ their expected return. Similarly, $\sigma_Y^{(n)}$ designates the volatility matrix of the state variables $Y_t^{(n)}$. All of these quantities are detailed in Appendix B. We solve the program $P_{ro}^{(n)}$; by directly applying the results of Liu (2007) but use a slightly different presentation:

Proposition 1.

The certainty equivalent of wealth, $CE_t^{(n)} \equiv u^{-1}(J_t^{(n)})$, can be computed as follows:

$$CE_t^{(n)} = W_t^{(n)} ce^{(n)}(t, Y^{(n)}), \quad (20.4a)$$

$$ce^{(n)}(t, Y^{(n)}) = \exp \left(CE_0^{(n)}(T_I - t) + CE_Y^{(n)}(T_I - t)' Y_t^{(n)} + \frac{1}{2} Y_t^{(n)'} CE_{YY}^{(n)}(T_I - t) Y_t^{(n)} \right), \quad (20.4b)$$

The optimal proportions can be decomposed as follows:

$$\pi_t^{(n)} = \frac{1}{\gamma} \pi_{Mvt}^{(n)} + \left[1 - \frac{1}{\gamma} \right] \pi_{CEht}^{(n)}, \quad (20.5a)$$

$$\pi_{Mvt}^{(n)} \equiv \Sigma_A^{(n)}(t)^{-1} [\mu_{At}^{(n)} - 1^{(n)} r_t], \quad (20.5b)$$

$$\pi_{CEht}^{(n)} \equiv -\Sigma_A^{(n)}(t)^{-1} \sigma_A^{(n)}(t)' \sigma_Y^{(n)} \frac{\partial_{Y^{(n)}} ce^{(n)}(t, Y^{(n)})}{ce^{(n)}(t, Y^{(n)})}, \quad (20.5c)$$

$$\frac{\partial_{Y^{(n)}} ce^{(n)}(t, Y^{(n)})}{ce^{(n)}(t, Y^{(n)})} = CE_Y^{(n)}(T_I - t) + CE_{YY}^{(n)}(T_I - t) Y_t^{(n)}, \quad (20.5d)$$

where the deterministic functions, $CE_0^{(n)}$, $CE_Y^{(n)}$, $CE_{YY}^{(n)}$, are given in Appendix B. $\partial_{Y^{(n)}}$ designates the gradient vector with respect to the state variables $Y^{(n)}$.

Proof. Available from the author upon request.

Proposition 1 confirms a well-known result: the optimal proportions are decomposed into a mean-variance efficient portfolio, $\pi_{Mvt}^{(n)}$, and an opportunity set hedging portfolio, $\pi_{CEht}^{(n)}$ (Merton 1971, 1973, Breeden 1979, 1984). $\pi_{Mvt}^{(n)}$ characterizes the momentum strategies because it only considers instantaneous asset movements: the excess return of assets over the risk-free rate, balanced out by the variance-covariance matrix of the same assets. It is nevertheless well known that the return of a futures price is a pure risk premium (Cox *et al.* 1981, Duffie and Stanton 1992). As a consequence, the excess return should be replaced by the return when a futures contract is considered.

This figure displays selected interest rates (time to maturities 1 month, 5 years and 20 years), [Figure 20.1a](#), as well as the (log) S&P 500 index, [Figure 20.1b](#), for weekly data ranging from 11 September 2009 to 14 March 2014.

$1^{(n)}$ is then a vector of ones and zeros, which is given in Appendix B. The weight that is invested in this speculative portfolio is the usual risk tolerance $1/\gamma$.

Traditionally, the hedging portfolio is decomposed into as many portfolios as the number of state variables: the well-known Merton-Breeden hedging terms, Merton (1971, 1973) and Breeden (1979, 1984). More recently, it has been proved that these terms relate to the dynamics of the short rate and of the market price of risks (e.g. Lioui and Poncet 2001, Munk and Sørensen 2007). In Proposition 1, we take a different approach and couch the hedging portfolio, $\pi_{CEht}^{(n)}$, in terms of the certainty equivalent of wealth. The interpretation of $\pi_{CEht}^{(n)}$ is straightforward. $\pi_{CEht}^{(n)}$ is indeed the opposite of the regression coefficients of $ce^{(n)}(t, Y^{(n)})$ on the traded assets. These regression coefficients involve the sensitivities of




FIGURE 20.1 Default-free yields and Equity time series.

$ce^{(n)}(t, Y^{(n)})$ to the state variables because the time-varying behaviour of $ce^{(n)}(t, Y^{(n)})$ stems from that of the state variables. Equation (20.5c) demonstrates that any unfavourable impact of the state variables on $ce^{(n)}(t, Y^{(n)})$ is offset by a favourable movement on the portfolio value, $W_t^{(n)}$. Therefore, this hedging procedure tends to smooth out the movements of the certainty equivalent $CE_t^{(n)}$, which is the product of $ce^{(n)}(t, Y^{(n)})$ and $W_t^{(n)}$.

This figure displays selected futures contracts (time to maturities 1 month and 24 months) for weekly data ranging from 11 September 2009 to 14 March 2014 for copper, gold and oil in Figure 20.2a, and c, respectively. Prices are in \$.

The weight invested in the portfolio, $\pi_{CEht}^{(n)}$, is naturally the complement to the unity of the weight invested in the mean-variance portfolio, $1 - 1/\gamma$. This weight is an increasing function of risk aversion. The more risk averse the investor, the more he/she will invest in this certainty equivalent hedge portfolio. This weight is positive for an investor who is more risk averse than the Bernoulli investor and negative otherwise. This decomposition is consistent with most empirical risk aversion studies (e.g. Meyer and




Figure 20.2 Futures on copper, gold and oil time series.

Meyer 2005), which find a risk aversion superior to one. Most theoretical allocation studies* also focus on a risk aversion that is superior to one. Indeed, an investor with a risk aversion below one will short the hedge portfolio, $\pi_{CEH}^{(n)}$ to invest more in the speculative portfolio. This type of investor behaviour appears unlikely in general.

20.3 Empirical Assessment of the Strategies

In this section, we implement the results of [Section 20.2](#) for three weakly economically related commodities: gold, copper and oil. Indeed, oil and gold are among the commodities that are most commonly used as investment vehicles as shown by the monthly reports of the single commodity indexes of the CFTC. Copper also sees significant use as an investment vehicle and is frequently studied in the academic literature (e.g. Schwartz 1997, Nielsen and Schwartz 2004, Casassus and Collin-Dufresne 2005, Tang 2012). Moreover, these three commodities are good representatives for the three main commodity sectors: energy (oil), industrial metal (copper) and precious metal (gold). [Section 20.3.1](#) is devoted to the presentation of the data. [Section 20.3.2](#) analyses the parameters and state variables implied by the data. [Section 20.3.3](#) discusses the results achieved by the implementation of optimal strategies.

20.3.1 Data Description

The data sample period goes from 11 September 2009, i.e. around one year after the collapse of Lehman Brothers, until 14 March 2014. We analyse weekly data. There is a total of 236 weeks—we collect data every Friday or on Thursday when the market is closed on the Friday. The default-free bond data are collected from the Federal Reserve and span times to maturity ranging from 1 month to 30 years. Selected default-free yields are displayed in [Figure 20.1\(a\)](#). The equity index is the S&P 500 and is plotted in [Figure 20.1\(b\)](#).

The futures contracts on copper, gold and oil are traded on the Chicago Mercantile Exchange and are identified by the product symbols, ‘HG’, ‘GC’ and ‘CL’, respectively. The equity index and the commodity futures contracts are obtained from Bloomberg. As in Schwartz (1997), we sort out futures contracts to obtain futures contracts with fixed time-to-maturities. The 1-month and 24-months futures contracts are displayed in [Figure 20.2\(a\)–\(c\)](#) for copper, gold and oil, respectively. The usual pattern persists: copper and oil alternate periods of Backwardation and Contango while gold is in Contango most of the time.

We obtain futures contracts ranging from 1 to 24, 48 and 60 months for copper, gold and oil, respectively. Descriptive statistics for these commodities are displayed in [Table 20.1](#).

Table 20.1 shows that on average, the copper term structure presents a hump shape increasing until the 12 months’ time-to-maturity contract, while gold clearly displays a Contango pattern. Oil displays a Backwardation pattern except for the first time-to-maturity. In addition, the standard deviation figures indicate that the Samuelson effect holds for copper and oil—except for the first time to maturity for copper and the last time to maturity for oil. However, the Samuelson effect does not hold for gold as predicted by Routledge *et al.* (2000). Indeed, gold is not much used in production and thus it displays a high level of inventory.

20.3.2 Parameters and State Variables Inference

Our factor model necessitates the estimation of 22 parameters for each commodity. Because of this high number of parameters, we follow Schwartz (1997) and choose to estimate the parameters of each market separately. We then estimate the correlation coefficients using the inferred state variables. For the bond

* See e.g. the various allocation studies cited by Munk and Sørensen (2007).

TABLE 20.1 Average Values for Futures on Copper, Gold and Oil

	Copper	Gold	Oil
F01	353.96 (41.60)	1433.08 (221.88)	90.78 (10.55)
F03	354.83 (41.91)	1434.12 (222.10)	91.83 (9.67)
F06	356.00 (41.81)	1437.09 (222.91)	92.05 (8.88)
F12	356.77 (41.11)	1442.48 (224.20)	91.67 (7.70)
F18	356.09 (39.82)	1449.63 (225.17)	90.81 (6.89)
F24	354.52 (38.54)	1451.75 (224.82)	90.00 (6.34)
F36	NA	1471.72 (221.61)	88.79 (5.92)
F48	NA	NA	88.27 (5.92)
F60	NA	NA	88.23 (6.21)

Notes: This table displays for selected futures contracts time to maturities ranging from 1 month to 60 months the average values for copper, gold and oil, for weekly data ranging from 11 September 2009 to March 14 2014. Standard deviations are in parenthesis.

and commodity markets, we choose to fit the principal components of their logarithm as in Casassus and Collin-Dufresne (2005).^{*} The (log) likelihood functions are displayed in Appendix C.[†]

For later reference, we denote by MV (Mean-Variance) and Dyn (Dynamic) any quantity that pertains to a constant and a dynamic market price of risk, respectively. This convenient notation is a slight misuse of language because a hedging term related to the interest rate is present when the market price of risk is constant (Munk and Sørensen 2004). However, the bond market is not our primary focus and this term is deterministic in our Gaussian framework (Munk and Sørensen 2004). Finally in order to save space, we provide only state variables inferred in the Dyn case: state variables stemming from the MV framework are very similar to the ones obtained in the Dyn setting and are available from the author upon request.

First, we provide the state variable[‡] and parameters inferred from the equity and bond markets' estimation.

Figure 20.3 shows that the pattern of the instantaneous risk-free rate is very similar to the patterns of the 5 years and 20 years default-free yield displayed in Figure 20.1, but differs from that of the one month yield displayed in the same figure. Indeed, our short rate is implied from the whole term structure of

^{*} In order to limit the number of parameters, we use a simplified version of their estimation and choose to fit perfectly the first principal component for the bond market and the two first principal components for the commodity markets. However, the first two principal components for the commodities represent more than 99% of the variance.

[†] As in most empirical studies similar to ours, see Hamilton and Wu (2012), we use the Matlab unconstrained optimization function, `fminunc`, to maximize our (log) likelihood function. When parameters are constrained to be positive, such as volatilities and speed of mean-reversion, we use an exponential-type transformation as suggested by Kim and Nelson (1999). Similarly, we use a tangent-type transformation for the correlation between the commodity spot price and its convenience yield.

[‡] We only provide the instantaneous risk-free as a state variable because the (log) equity price is observable and is displayed in Figure 20.1(b).




Figure 20.3 Implied short rate from the bond market.

interest rates. In addition, the short rate is lower than the 5 years default-free yield, which confirms a generally increasing term structure of interest rates.

Table 20.2 shows that the volatilities of the equity return and the short rate are in line with usual market values. The long-term mean of the short rate is very small and confirms the low level of interest rates during the period of study. The equity risk premium for the MV case, β_{s0} , is positive as shown by the increasing pattern of the S&P 500 in Figure 20.1. β_{r0} is negative in the MV case, yet very low in absolute value. It indicates a positive but very low-risk premium for the bond market. The speed of mean-reversion of the short rate is in line with the study of Casassus and Collin-Dufresne (2005) for the MV case, but higher in the Dyn setting. The risk-neutral speed of mean-reversion of the short rate, $\kappa_r = \kappa_r + \beta_{rr} = 8.65\%$, is also higher than in the study of Casassus and Collin-Dufresne (2005). This difference may stem from a shorter period of study of our analysis. The correlation coefficient between the short rate and the equity index is computed using the implied short rate and is the same in both the MV and the Dyn cases: $\rho_{sr} = 58.90\%$. It is positive as in Brennan and Xia (2000).

A (log) likelihood ratio test available from the author upon request proves that the Dyn model does not statistically improve the estimation. However, a closer inspection of Table 20.2 shows that, in general, the Dyn model provides more significant parameters than the MV estimation i.e. with lower standard errors. We obtain a negative β_{rr} as in Casassus and Collin-Dufresne (2005). β_{ss} is negative with a low standard error, which is in line with the mean-reverting pattern of equity outlined in the study of Wachter (2002).

Figure 20.4 shows the implied (log) spot commodity price and its convenience yield for our period of study. As in Schwartz (1997), we add a constant value to the convenience yield to compare it with the (log) spot price on the same graph. Figure 20.4 suggests intuitions that the three convenience yields and

TABLE 20.2 Equity and Bond Markets Implied Parameters

(a) Equity parameters

MV

	σ_s	β_{s0}
Parameters	15,58%	11,71%
Std. err.	4,61%	7,33%
Log likelihood	-144,80	

Dyn

	σ_s	β_{s0}	β_{ss}
Parameters	15,55%	225,23%	-29,43%
Std. err.	4,61%	7,31%	1,02%
Log likelihood	-144,38		

(b) Bond parameters

MV

	κ_r	θ_r	σ_r	β_{r0}
Parameters	32,88%	0,60%	2,48%	-0,68%
Std. err.	5,44%	0,82%	4,61%	0,35%
Log likelihood	-115,59			

Dyn

	κ_r	θ_r	σ_r	β_{r0}	β_{rr}
Parameters	82,35%	0,26%	1,01%	-0,23%	-73,70%
Std. err.	1,01%	0,07%	4,60%	0,07%	0,82%
Log likelihood	-115,15				

Note: This table displays the parameters and their standard error (Std. err.) for the equity and bond markets with the (log) likelihood given in subappendixes (Ca, b) for the MV and Dyn cases.




FIGURE 20.4 Implied (log) spot price and convenience yield for copper, gold and oil.

TABLE 20.3 Correlation Coefficients Implied from State Variables

	<i>MV</i>			<i>Dyn</i>			
	ρ_{sx}	ρ_{ss}	ρ_{rx}	ρ_{rs}	ρ_{sx}	ρ_{ss}	ρ_{rx}
<i>Gold</i>							
57.33%	36.12%	43.18%	45.23%	57.33%	36.12%	43.20%	45.23%
<i>Copper</i>							
24.01%	44.04%	-4.46%	60.79%	24.60%	44.04%	-3.65%	60.79%
<i>Oil</i>							
53.21%	28.67%	40.82%	35.28%	54.02%	28.67%	40.64%	35.28%

Note: This table displays the correlation coefficients for the (log) equity price, instantaneous risk-free rate, (log) spot commodity and convenience yield (copper, gold and oil) for the MV and Dyn cases.

(log) spot prices exhibit a mean-reverting pattern; the convenience yields and (log) spot prices of the copper and oil commodities are highly positively correlated; the convenience yield of gold is low and exhibits a low volatility.

This figure displays the instantaneous risk-free rate implied from the bond market in the Dyn case for weekly data ranging from 11 September 2009 to 14 March 2014.

Table 20.3 confirms a tendency outlined in the study of Rouwenhorst and Tang (2012): since the end of the 2000s, the correlation coefficient between equity and commodities is relatively high. We outline in our study the same phenomena between the short rate and copper and oil—the correlation between the short rate and gold is close to zero. This level of correlation is quite different from the parameters obtained with market data from the 1990s where the correlation coefficient between the risk-free rate and the spot commodity and its convenience yield is close to zero (Schwartz 1997, Casassus and Collin-Dufresne 2005).

Table 20.4 shows that the parameters obtained for the convenience yields and the volatilities of the spot commodities are similar to those obtained in previous studies (Schwartz 1997, Casassus and Collin-Dufresne 2005). A notable exception is the correlation coefficient between the spot commodity and its convenience yield for gold which is close to zero in our study, but it is usually slightly positive (Schwartz 1997, Casassus and Collin-Dufresne 2005). Table 20.4 also reports that the parameters obtained in the MV and in the Dyn cases are similar. Finally, the comparison between the results of Table 20.4 and the

TABLE 20.4 Parameters for the Copper, Gold and Oil Markets

(a) Copper

MV

	β_{x0}	σ_x	σ_δ	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	10.66%	27.97%	20.54%	256.16%	7.48%	11.24%	67.52%
Std. err.	9.71%	3.84%	3.64%	2.71%	2.78%	0.52%	19.22%
Log Likelihood	1291.71						

Dyn

	β_{x0}	$\beta_{x\delta}$	β_{xx}	σ_x	σ_δ	$\kappa_{\delta Q}$	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	1078.33%	188.03%	-183.69%	27.64%	20.82%	258.48%	229.67%	7.25%	11.34%	67.33%
Std. err.	9.61%	72.49%	1.63%	3.83%	3.56%	2.71%	30.66%	3.15%	0.43%	19.06%
Log Likelihood	1294.04									

(b) Gold

MV

	β_{x0}	σ_x	σ_δ	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	7.50%	17.96%	2.10%	98.73%	1.07%	5.59%	-0.15%
Std. err.	8.45%	4.61%	4.46%	4.92%	1.00%	0.48%	10.23%
Log Likelihood	1661.69						

Dyn

	β_{x0}	$\beta_{x\delta}$	β_{xx}	σ_x	σ_δ	$\kappa_{\delta Q}$	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	948.08%	-55.54%	-129.24%	17.79%	3.17%	146.06%	107.54%	1.35%	6.39%	-0.33%
Std. err.	8.37%	371.33%	1.15%	4.61%	4.30%	3.90%	68.64%	1.39%	0.36%	10.22%
Log likelihood	1663.96									

(c) Oil

MV

	β_{x0}	σ_x	σ_δ	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	11.86%	29.56%	25.66%	110.06%	8.09%	-9.24%	80.28%
Std. err.	8.29%	3.32%	3.33%	2.87%	6.54%	1.94%	30.50%
Log Likelihood	822.07						

Dyn

	β_{x0}	$\beta_{x\delta}$	β_{xx}	σ_x	σ_δ	$\kappa_{\delta Q}$	κ_δ	θ_δ	$\theta_{\delta Q}$	$\rho_{x\delta}$
Parameters	1503.71%	313.79%	-368.91%	27.79%	20.58%	92.94%	112.73%	9.20%	-11.60%	78.31%
Std. err.	8.13%	61.81%	2.00%	3.38%	3.40%	2.97%	37.19%	5.34%	1.01%	27.75%
Log Likelihood	826.74									

Note: This table displays the parameters and their standard error (Std. err.) for the dynamics of the (log) spot commodity and convenience yield (copper, gold and oil) for the MV and Dyn cases.

parameters obtained by Schwartz (1997) and Casassus and Collin-Dufresne (2005) highlights the fact that the volatility of gold (oil) is higher (lower) in the 2010s than in the 1990s.

This figure displays the (log) spot price and its convenience yield for copper, gold and oil for weekly data ranging from 11 September 2009 to 14 March 2014.

The likelihood achieved in the Dyn case is higher than in the MV case, but a log likelihood ratio test available from the author upon request shows that this difference is significant only for oil. Finally, Table 20.4 confirms the findings of Casassus and Collin-Dufresne (2005): β_{xx} is negative for the three commodities with very low standard error, while the standard error of β_{x6} is relatively high.

20.3.3 Optimal Strategies

In this paragraph, we compare different types of strategies. First, we compare the benefit of investing in the commodity in addition to an investment in bonds and stocks. For this commodity investment, two subcases are considered. A subcase where the investor invests in a distant futures contract in addition to the spot commodity: we abbreviate this strategy as Spr in reference to the well-known calendar spread strategy. A subcase where the investor restricts himself/herself to the spot commodity: we abbreviate this strategy by She as short end in reference to the short end of the term structure represented by the spot commodity price.

We present our results for a distant futures contract and a long-term bond with time to maturities 1 year and 20 years, respectively. The sorting out of our futures contracts as in Schwartz (1997) ensures that the one-year futures contract is always available to the investor for the three commodities. We report our results for an investment that covers our estimation period as in Brennan *et al.* (1997): a weekly change of positions starting from 11 September 2009 and ending on 14 March 2014. Further, we consider an investor with an additional investment horizon of 5 years after our last sample date, i.e. 14 March 2014. We also computed our strategies for different remaining investment horizons (0, 1 year and 10 years), but our results are not affected by the investment horizon. We provide the two first empirical moments of the strategies: mean and standard deviation in order to save space.* These two statistics are a convenient way to sum up the impact of state variables on the investment. Finally, we exhibit our results for four levels of risk aversion as is usual in the asset allocation literature, (e.g. Brennan and Xia 2000): $\gamma = 1, 3, 5, 9$.

First, Table 20.5 confirms well-known results, the state variables and investment horizon do not impact stock, spot commodity and futures contract proportions when the market price of risk is constant (see e.g. Munk 2013). In the same vein, results available from the author upon request prove that the standard deviation of the bond proportion in the MV case only stems from the investment horizon in our Gaussian setting (Munk and Sørensen 2004). Second, Table 20.5 conforms to intuition as far as risk aversion is concerned: optimal proportions are in general a decreasing function of risk aversion. Indeed, Table 20.5 shows that optimal proportions are in general very high for aggressive strategies ($\gamma = 1$).

Table 20.5 confirms that a dynamic market price of risk results in proportions with high standard deviations that may be hard to implement. Regarding the Spr strategy, Table 20.5 demonstrates that on average, the optimal strategy for oil is a calendar spread while an optimal investment for copper and gold results in a reverse calendar spread, i.e. off-setting positions. As far as copper and oil are concerned, the calendar spread is never one-to-one. Regarding gold, the Spr strategy leads to extreme positions in the term structure of futures price that are almost equal to each other in absolute value.

We now turn to the implementation of optimal strategies. The spot commodity is not available for investing and is replaced with the futures contract with the shortest time to maturity. Computations available from the author upon request leads to the following equation for the return on wealth:

$$\begin{aligned} \frac{\Delta W_t}{W_t} = & r_t \Delta t + \pi_{S_t} \left(\frac{\Delta S_t}{S_t} - r_t \Delta t \right) + \pi_{P_t} \left(\frac{\Delta P_t}{P_t} - r_t \Delta t \right) \\ & + \pi_{F_{1t}} \frac{\Delta F_{1t}}{F_{1t}} + \pi_{F_{2t}} \frac{\Delta F_{2t}}{F_{2t}}, \end{aligned} \quad (20.6)$$

* All values are available from the author upon request.

TABLE 20.5 Optimal Proportions for Equity, Bond, the Commodity and the Distant Futures Contract

		$\gamma = 1$				$\gamma = 3$				$\gamma = 5$				$\gamma = 9$			
		πS	πB	πX	πF	πS	πB	πX	πF	πS	πB	πX	πF	πS	πB	πX	πF
(a) Bond/stock																	
Bond/stock MV	Mean	4.83	3.66	NA	NA	1.61	2.81	NA	NA	0.97	2.64	NA	NA	0.54	2.53	NA	NA
	Std	0.00	0.00	NA	NA	0.00	0.29	NA	NA	0.00	0.35	NA	NA	0.00	0.38	NA	NA
Bond/Stock Dyn	Mean	11.85	17.53	NA	NA	7.07	11.96	NA	NA	5.30	9.67	NA	NA	3.60	7.29	NA	NA
	Std	5.44	11.14	NA	NA	2.88	6.02	NA	NA	2.15	4.47	NA	NA	1.49	3.15	NA	NA
(b) Bond/stock/copper																	
Bond/stock/copper Spr MV	Mean	7.56	14.95	-11.24	13.96	2.52	5.58	-3.75	4.65	1.51	3.71	-2.25	2.79	0.84	2.46	-1.25	1.55
	Std	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.04	0.00	0.00
Bond/stock/copper Spr Dyn	Mean	10.26	16.39	-9.79	12.64	5.68	11.11	-3.18	6.18	4.13	9.09	-2.00	5.08	2.70	6.88	-1.23	4.34
	Std	5.97	12.16	7.12	12.28	2.97	6.54	2.94	6.23	2.14	4.82	1.95	4.49	1.44	3.37	1.20	3.06
Bond/stock/copper She MV	Mean	8.73	14.76	0.29	NA	2.91	5.52	0.10	NA	1.75	3.67	0.06	NA	0.97	2.44	0.03	NA
	Std	0.00	0.00	0.00	NA	0.00	0.03	0.00	NA	0.00	0.03	0.00	NA	0.00	0.04	0.00	NA
Bond/stock/copper She Dyn	Mean	11.32	17.78	0.61	NA	6.44	12.51	0.68	NA	4.77	10.33	0.65	NA	3.24	7.95	0.55	NA
	Std	5.24	12.16	4.44	NA	2.70	6.53	1.98	NA	1.99	4.83	1.32	NA	1.36	3.39	0.80	NA
(c) Bond/stock/gold																	
Bond/stock/gold Spr MV	Mean	11.67	61.20	-250.52	250.56	3.89	21.00	-83.51	83.52	2.33	12.96	-50.10	50.11	1.30	7.60	-27.84	27.84
	Std	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.04	0.00	0.00
Bond/stock/gold Spr Dyn	Mean	14.79	19.33	-207.69	207.41	8.77	7.34	-58.45	58.11	6.42	5.63	-34.28	34.22	4.23	4.15	-18.85	19.13
	Std	4.55	9.28	57.64	53.27	2.58	5.76	17.66	15.11	1.95	4.33	10.40	8.45	1.36	3.03	5.66	4.26
Bond/stock/gold She MV	Mean	8.89	14.47	0.20	NA	2.96	5.42	0.07	NA	1.78	3.61	0.04	NA	0.99	2.41	0.02	NA
	Std	0.00	0.00	0.00	NA	0.00	0.03	0.00	NA	0.00	0.03	0.00	NA	0.00	0.04	0.00	NA
Bond/stock/gold She Dyn	Mean	11.83	17.51	0.04	NA	7.16	12.02	-0.24	NA	5.34	9.70	-0.16	NA	3.61	7.30	-0.03	NA
	Std	4.77	9.64	5.99	NA	2.65	5.32	3.50	NA	1.99	3.99	2.72	NA	1.37	2.83	2.04	NA
(d) Bond/stock/oil																	
Bond/stock/oil Spr MV	Mean	12.15	15.36	7.47	-13.06	4.05	5.72	2.49	-4.35	2.43	3.79	1.49	-2.61	1.35	2.51	0.83	-1.45
	Std	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.04	0.00	0.00
Bond/stock/oil Spr Dyn	Mean	15.31	20.42	12.10	-18.68	9.11	14.24	6.51	-12.12	6.72	11.52	4.92	-9.81	4.51	8.76	3.67	-7.84
	Std	9.24	11.62	5.42	17.00	4.16	6.35	3.11	9.71	2.94	4.74	2.42	7.47	1.93	3.36	1.82	5.52
Bond/stock/oil She MV	Mean	8.56	14.90	0.51	NA	2.85	5.57	0.17	NA	1.71	3.70	0.10	NA	0.95	2.46	0.06	NA
	Std	0.00	0.00	0.00	NA	0.00	0.03	0.00	NA	0.00	0.03	0.00	NA	0.00	0.04	0.00	NA
Bond/stock/oil She Dyn	Mean	10.84	18.01	1.23	NA	6.49	12.62	0.62	NA	4.91	10.30	0.43	NA	3.39	7.81	0.27	NA
	Std	6.01	11.75	4.79	NA	2.87	6.32	1.92	NA	2.10	4.68	1.21	NA	1.44	3.28	0.69	NA

Notes: This table gives optimal proportions given by Equations (20.5a)–(20.5d) for the period starting from 11 September 2009 and finishing 14 March 2014, a remaining investment horizon of 5 years and risk aversions $\gamma = 1, 3, 5$ and 9. The definitions of asset proportions are given in Appendix B.

where F_{1t} stands for the shortest time to maturity futures contract and F_{2t} is the distant futures contract—one year time to maturity in our case. We use Equation (20.6) to compute the excess returns of our strategies as well as their Sharpe ratio. As for the optimal portfolio proportions, we display the mean values of these excess returns as well as their standard deviations. Results are presented in Table 20.6(a).

To further analyse the Sharpe ratios, we use the statistic developed by Jobson and Korkie (1981), (see also Basu and Miffre 2013), which shows that the difference between two Sharpe ratios is distributed as a z statistic:

$$z = \frac{\lambda_1 - \lambda_2}{\sqrt{\frac{1}{T} \left[2 - 2\rho_{1,2} + \frac{1}{2}\lambda_1^2 + \frac{1}{2}\lambda_2^2 - \lambda_1\lambda_2\rho_{1,2}^2 \right]}}, \quad (20.7)$$

where λ_1, λ_2 are the Sharpe ratios of two strategies, $\rho_{1,2}$, their correlation coefficient and T is the number of observations. The values of the z statistics are given in Table 20.6(b).

Table 20.6(a) shows that the mean excess returns and their standard deviations are decreasing functions of risk aversion. The Sharpe ratio is in general a decreasing function of risk aversion except when a gold (reverse) calendar spread is considered. In addition, Table 20.6(a) shows that the volatility of the strategy is higher when a dynamic market price of risk is considered except for the gold calendar spread.

Table 20.6(a) confirms that the Sharpe ratios of the Dyn strategies are higher than those of the MV strategies. Table 20.6(b) shows, nevertheless, that the differences in the Sharpe ratios are not statistically significant. In addition, Table 20.6(a) proves that the calendar spread strategies do not provide a higher Sharpe ratio over the simple Bond/Stock strategies. However, Table 20.6(a) also demonstrates that investing only in the short end of the commodity term structure provides a better Sharpe ratio than investing in both ends of the term structure. This result may be due to the extreme opposite positions taken by investors on the term structure of futures prices:* the results achieved for the calendar spreads are very sensitive to the asset returns of the two futures contracts when they are highly correlated. Indeed, computations available upon request show that the correlation between the 1 month and the 24 months futures contracts equals 88, 99 and 85% for copper, gold and oil, respectively. The difficulties linked to an investment in both ends of the term structure will be further highlighted by the results provided in Table 20.7. Finally, Table 20.6(a) shows that an investment in the spot commodity in addition to an investment in bonds and stocks provides a better Sharpe ratio when the Dyn case is considered.†

Table 20.7 reports the wealth obtained for a 100\$ initial investment and compounding the return given by Equation (20.6). First, Table 20.7 shows that a Dyn strategy achieves a higher wealth than MV in all cases for a bond-stock mix together with investment in the spot commodity. Moreover, investing in the spot commodity in addition to bond and equity also achieves a better terminal wealth. Second, Table 20.7 highlights the bad results achieved by the spread strategies. Indeed, the extreme calendar spread positions on the term structure of commodity futures prices lead to a terminal wealth that is lower than that achieved by a simple bond-stock mix except when the Dyn strategy is considered and the investor is conservative: $\gamma \geq 5$ for gold and oil and $\gamma \geq 3$ for copper. For aggressive strategies ($\gamma = 1$), the terminal wealth can even turn out to be negative as far as gold and oil are concerned.

20.4 Concluding Remarks

This article focuses on commodity investment in light of the financialization of commodities (Rouwenhorst and Tang 2012) in a dynamic framework that is compatible with the theory of storage (Working 1948, 1949). We analyse the benefit of investing in commodities in addition to a simple bond-stock mix strategy. We consider two types of investment: an investment only in the spot commodity and an investment in both the spot commodity and a distant commodity futures contract. Our strategies

* We thank an anonymous referee for this suggestion.

† This is the important result referred to in the abstract.

TABLE 20.6(A) Excess Return and Sharpe Ratio of Optimal Strategies

	Bond/stock MV				Bond/stock Dyn											
	1	3	5	9	1	3	5	9								
γ																
Mean	1.46	0.58	0.41	0.29	5.18	3.14	2.39	1.67								
Std	8.20	3.26	2.99	3.11	24.54	15.01	11.64	8.45								
Sharpe Ratio	17.85	17.91	13.61	9.34	21.12	20.93	20.49	19.73								
	Bond/stock/copper Spr MV				Bond/stock/copper Spr Dyn				Bond/stock/copper She MV				Bond/stock/copper She Dyn			
	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
γ																
Mean	2.78	0.96	0.59	0.35	7.06	4.27	3.31	2.42	3.08	1.06	0.65	0.38	6.70	3.76	2.81	1.94
Std	19.42	6.89	4.44	2.86	32.55	20.50	17.01	14.03	17.72	6.40	4.19	2.76	28.25	16.36	12.68	9.29
Sharpe Ratio	14.33	13.92	13.40	12.29	21.70	20.83	19.44	17.25	17.37	16.52	15.62	13.92	23.70	23.00	22.13	20.92
	Bond/stock/gold Spr MV				Bond/stock/gold Spr Dyn				Bond/stock/gold She MV				Bond/stock/gold She Dyn			
	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
γ																
Mean	5.66	1.92	1.17	0.67	7.50	4.41	3.35	2.39	3.11	1.07	0.66	0.39	7.13	4.25	3.24	2.31
Std	91.30	31.15	19.12	11.12	58.95	23.26	16.10	10.59	17.63	6.37	4.16	2.75	28.60	17.40	13.53	9.91
Sharpe Ratio	6.20	6.16	6.12	6.04	12.72	18.95	20.82	22.61	17.66	16.80	15.89	14.15	24.92	24.40	23.92	23.28
	Bond/stock/oil Spr MV				Bond/stock/oil Spr Dyn				Bond/stock/oil She MV				Bond/stock/oil She Dyn			
	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
γ																
Mean	3.98	1.36	0.83	0.48	10.83	6.37	4.79	3.36	3.13	1.07	0.66	0.39	7.45	4.13	3.03	2.06
Std	23.30	8.24	5.27	3.33	51.43	32.02	25.76	20.30	17.90	6.47	4.23	2.79	31.00	17.07	12.94	9.27
Sharpe Ratio	17.07	16.48	15.83	14.54	21.06	19.90	18.58	16.57	17.45	16.59	15.69	13.99	24.03	24.18	23.45	22.22

Note: This table gives the excess return, mean and standard deviation (Std) and the Sharpe ratio of optimal strategies given by equations for the period starting from 11 September 2009 and finishing 14 March 2014, a remaining investment horizon of 5 years and risk aversions $\gamma = 1, 3, 5$ and 9.

TABLE 20.6(B) z Statistic for the Comparison Sharpe Ratio of Optimal Strategies

γ	Copper				Gold				Oil			
	1	3	5	9	1	3	5	9	1	3	5	9
	Bond/stock/commodity Spr MV vs. Bond/Stock MV											
zStat	0.75	1.58	0.07	0.84	1.27	1.99	1.70	0.89	0.12	0.35	0.57	1.38
	Bond/stock/commodity She MV vs. bond/stock MV											
zStat	0.11	1.79	1.01	1.68	0.04	1.54	1.12	1.75	0.09	1.44	1.04	1.71
	Bond/stock/commodity Spr MV vs. bond/stock/commodity She MV											
zStat	1.33	1.21	1.11	0.95	1.82	1.85	1.84	1.75	0.08	0.03	0.04	0.16
	Bond/stock/commodity Spr Dyn vs. bond/stock Dyn											
zStat	0.12	0.02	0.20	0.42	1.18	0.32	0.05	0.47	0.01	0.16	0.29	0.47
	Bond/stock/commodity Spr Dyn vs. bond/stock/commodity She Dyn											
zStat	1.02	0.74	0.75	0.80	1.82	1.02	0.62	0.14	0.67	0.82	0.88	0.96
	Bond/stock/commodity She Dyn vs. bond/stock Dyn											
zStat	0.72	0.75	0.66	0.53	1.00	0.93	0.92	0.94	0.67	1.04	1.14	1.19
	Bond/stock/commodity Spr Dyn vs. bond/stock/commodity Spr MV											
zStat	1.34	1.36	1.15	0.83	1.11	1.63	1.82	2.02	1.11	1.63	1.82	2.02
	Bond/stock/commodity She Dyn vs. Bond/stock/commodity She MV											
zStat	1.25	1.58	1.72	1.95	1.43	1.59	1.70	1.92	1.20	1.77	2.01	2.32
	Bond/stock Dyn vs. Bond/Stock MV											
zStat	0.79	1.02	1.68	2.16								

Note: This table gives the excess return, mean and standard deviation (Std), and the Sharpe ratio of optimal strategies given by equations for the period starting from 11 September 2009 and finishing 14 March 2014, a remaining investment horizon of 5 years and risk aversions $\gamma = 1, 3, 5$ and 9.

are compared with the bond-stock mix benchmark (Brennan and Xia 2000). Our study is carried over three weakly economically linked commodities: gold, copper and oil. Each strategy is analysed for two embedded cases: a constant and a dynamic market price of risk.

We find that investing in both ends of the commodity term structure gives poor performance: the extreme calendar spread positions result in wealth, that is, in general, lower than that of a simple bond-stock mix strategy. While the Sharpe ratio analysis leads to mixed results, the analysis of the terminal wealth, nevertheless, clearly indicates that the addition of a spot commodity improves the investment. Finally, a dynamic market price of risk greatly improves the terminal value of the strategies. Our results thus extend the findings of Brennan and Schwartz (1998) to commodities to find an opposite conclusion: (distant) futures contracts on commodities do not in general improve the situation relative to investment in commodities themselves.*

* A referee has pointed out that an attempt to replicate the original Brennan–Schwartz results yielded a similar conclusion to that presented here, i.e. optimally investing in long-dated futures on stocks and bonds decreased returns.

TABLE 20.7 Wealth of Optimal Strategies for an Initial Endowment of 100\$

Bond/stock MV					Bond/stock Dyn											
γ	1	3	5	9	1	3	5	9								
Last	1525	381	257	194	9612	10,490	5360	2165								
Mean	532	233	185	158	2489	3589	2121	1014								
Std	404	97	61	43	2859	3537	1889	775								
Bond/stock/copper Spr MV					Bond/stock/copper Spr Dyn				Bond/stock/copper She MV				Bond/stock/copper She Dyn			
γ	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
Last	384	589	351	228	1020	16,038	8253	3114	2260	801	413	247	28,841	27,440	10,636	3455
Mean	2136	470	277	189	2252	19,457	9695	3658	2409	444	263	183	13,323	13,083	5453	1926
Std	2513	280	121	57	3284	21,263	9498	3106	2905	278	120	57	19,144	14,189	5239	1589
Bond/stock/gold Spr MV					Bond/stock/gold Spr Dyn				Bond/stock/gold She MV				Bond/stock/gold She Dyn			
γ	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
Last	0	-2	-18	111	-384,635	-47,867	7224	7236	2593	828	421	250	74,599	60,449	23,512	7197
Mean	2	108	223	220	-62,108	-3326	3091	2349	2490	447	264	183	11,760	13,740	6512	2504
Std	14	131	242	147	153,608	13,549	2857	2029	3009	282	122	58	16,838	15,950	6701	2198
Bond/stock/oil Spr MV					Bond/stock/oil Spr Dyn				Bond/stock/oil She MV				Bond/stock/oil She Dyn			
γ	1	3	5	9	1	3	5	9	1	3	5	9	1	3	5	9
Last	1444	1199	561	301	-1080	-47,905	6713	4975	2227	821	420	250	23,171	48,307	16,608	4545
Mean	785	449	270	186	-93	-10,174	1965	1663	2411	452	266	184	5563	15,115	6090	1998
Std	1141	352	153	70	653	14,897	1957	1530	2911	285	123	58	7296	16,145	5825	1642

Notes: This table gives the wealth obtained for an initial investment of 100\$: Last refers to the terminal value, Mean to the mean value and Std to the standard variation. The period under consideration starts from 11 September 2009 and finishes 14 March 2014, with a remaining investment horizon of 5 years and risk aversions $\gamma = 1, 3, 5$ and 9.

References

- Basak, S. and Pavlova, A., A model of financialization of commodities. Working Paper, London Business School, 2013.
- Basu, D. and Miffre, J., The performance of simple dynamic commodity strategies. *J. Altern. Invest.*, 2013, **16**, 9–18.
- Bessembinder, H., Systematic risk, hedging pressure and risk premiums in futures markets. *Rev. Financ. Stud.*, 1992, **5**, 637–667.
- Bessembinder, H. and Chan, K., Time-varying risk premia and forecastable returns in futures markets. *J. Financ. Econ.*, 1992, **32**, 169–193.
- Black, F., The pricing of commodity contracts. *J. Financ. Econ.*, 1976, **3**, 167–179.
- Brennan, M.J. and Schwartz, E., The use of treasury bill futures in strategic asset allocation programs. In *Worldwide Asset and Liability Modeling* edited by W.T. Ziemba and J.M. Mulvey, pp. 205–230, 1998 (Cambridge University Press: Cambridge).
- Brennan, M.J., Schwartz, E.S. and Lagnado, R., Strategic asset allocation. *J. Econ. Dyn. Control*, 1997, **21**, 1377–1403.
- Brennan, M.J. and Xia, Y., Stochastic interest rates and the bondstock mix. *Rev. Finance*, 2000, **4**, 197–210.
- Breedon, D., An intertemporal asset pricing model with stochastic consumption and investment opportunities. *J. Financ. Econ.*, 1979, **7**, 263–296.
- Breedon, D., Futures markets and commodity options—Hedging and optimality in incomplete markets. *J. Econ. Theory*, 1984, **32**, 275–300.
- Casassus, J.P. and Collin-Dufresne, J.-P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Cheng, I. and Xiong, W., The financialization of commodity markets. Working Paper, Dartmouth College and Princeton University, 2013.
- Cox, J., Ingersoll, J.E. and Ross, S., The relation between forward and futures prices. *J. Financ. Econ.*, 1981, **9**, 321–346.
- Dai, R., Commodities in dynamic asset allocation: Implications of mean reverting commodity prices. Working Paper, Tilburg University, 2009.
- Dai, Q. and Singleton, K., Expectation puzzle, time-varying risk premia and affine models of the term structure. *J. Financ. Econ.*, 2002, **63**, 415–441.
- de Roon, F., Nijman, T. and Veld, C., Hedging pressures effects in futures markets. *J. Finance*, 2000, **55**, 1437–1456.
- Duffee, G., Term premia and interest rate forecasts in affine models. *J. Finance*, 2002, **57**, 405–443.
- Duffie, D. and Stanton, R., Pricing continuously resettled contingent claim. *J. Econ. Dyn. Control*, 1992, **16**, 561–573.
- Erb, C. and Harvey, C., The strategic and tactical value of commodity futures. *Financ. Anal. J.*, 2006, **62**, 69–97.
- Fuertes, A.M., Miffre, J. and Rallis, G., Tactical allocation in commodity futures markets: Combining momentum and term structure signals. *J. Banking Finance*, 2010, **34**, 2530–2548.
- Fung, W. and Hsieh, D., Empirical characteristics of dynamic trading strategies: The case of hedge funds. *Rev. Financ. Stud.*, 1997, **10**, 275–302.
- Gibson, R. and Schwartz, E.S., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Gorton, G., Hayashi, F. and Rouwenhorst, K., The fundamentals of commodity futures returns. Working Paper, Yale University, 2007.
- Hamilton, J. and Wu, J., Identification and estimation of Gaussian affine term structure models. Working Paper, University of California and University of Chicago, 2012.
- Hong, H., Stochastic convenience yield, optimal hedging and the term structure of open interest and futures prices. Working Paper, Graduate School of Business, Stanford University, 2001.

- Hong, H. and Yogo, M., Commodity asset interest and asset returns predictability. Working Paper, Princeton University, 2010.
- Jobson, J.D. and Korkie, B., Performance hypothesis testing with the Sharpe and Treynor measures. *J. Finance*, 1981, **36**, 889–908.
- Kim, C.-J. and Nelson, C., *State-space Models with Regime Switching: Classical and Gibbs Sampling Approaches with Applications*, 1999 (Massachusetts Institute of Technology Press: Cambridge).
- Lioui, A. and Poncet, P., On the Optimal Portfolio Choice under Stochastic Interest Rates. *J. Econ. Dyn. Control*, 2001, **25**, 1841–1865.
- Liu, J., Portfolio selection in stochastic environments. *Rev. Financ. Stud.*, 2007, **20**, 1–39.
- Marshall, B., Cahan, R. and Cahan, J., Can commodity futures be profitably traded with quantitative market timing strategies? *J. Banking Finance*, 2008, **32**, 1810–1819.
- Munk, C., Dynamic asset allocation. Incomplete Lecture Notes, 2013 (Copenhagen Business School).
- Munk, C. and Sørensen, C., Optimal consumption and investment strategies with stochastic interest rates. *J. Banking Finance*, 2004, **28**, 1987–2013.
- Merton, R., Lifetime portfolio selection under uncertainty: The continuous-time case. *Rev. Econ. Stat.*, 1969, **51**, 247–257.
- Merton, R., Optimum consumption and portfolio rules in a continuous-time model. *J. Econ. Theory*, 1971, **3**, 373–413.
- Merton, R., An intertemporal capital asset pricing model. *Econometrica*, 1973, **41**, 867–887.
- Meyer, D. and Meyer, J., Relative risk aversion: What do we know? *J. Risk Uncertainty*, 2005, **31**(3), 243–262.
- Miffre, J., Backwardation is normal. *J. Futures Markets*, 2000, **20**, 803–821.
- Miffre, J. and Rallis, G., Momentum strategies in commodity futures markets. *J. Banking Finance*, 2007, **31**, 1863–1886.
- Munk, C. and Sørensen, C., Optimal real consumption and investment strategies in dynamic stochastic economies. In *Stochastic Economic Dynamics* edited by B.S. Jensen and T. Palokangas, pp. 271–316, 2007 (Copenhagen Business School Press: Copenhagen).
- Nielsen, M. and Schwartz, E., Theory of storage and the pricing of commodity contingent claims. *Rev. Derivatives Res.*, 2004, **7**, 5–24.
- Routledge, B., Seppi, D. and Spatt, C., Equilibrium forward curves for commodities. *J. Finance*, 2000, **55**, 1297–1338.
- Rouwenhorst, G. and Tang, K., Commodity investing. *Annu. Rev. Financ. Econ.*, 2012, **4**, 447–467.
- Scherer, B. and He, L., The diversification benefits of commodity futures indexes: A mean-variance spanning test. In *The Handbook of Commodity Investing* edited by F. Fabozzi, F. Füss and D. Kaiser, pp. 241–265, 2008 (Wiley: Hoboken, NJ).
- Schwartz, E., The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Singelton, K., Investor flows and the 2008 boom/bust in oil prices. Working Paper, Stanford University, 2012.
- Su, Y. and Keung, M.L.C., Strategic asset allocation and intertemporal hedging demands: With commodities as an asset class. Working Paper, Suffolk University, 2010.
- Tang, K., Time-varying long-run mean of commodity prices and the modeling of futures term structures. *Quant. Finance*, 2012, **12**, 781–790.
- Tang, K. and Xiong, W., Index investment and the financialization of commodities. *Financ. Anal. J.*, 2012, **68**, 54–74.
- Vasicek, O., An equilibrium characterization of the term structure. *J. Financ. Econ.*, 1977, **5**, 177–188.
- Wachter, J., Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. *J. Financ. Quant. Anal.*, 2002, **37**, 63–91.
- Willenbrock, S., Diversification return, portfolio rebalancing, and the commodity return puzzle. *Financ. Anal. J.*, 2011, **67**, 42–49.

Working, H., Theory of the inverse carrying charge in futures markets. *J. Farm. Econ.*, 1948, **30**, 1–28.

Working, H., The theory of the price of storage. *Amer. Econ. Rev.*, 1949, **39**, 1254–1262.

Appendix A: The Volatilities of the Zero-Coupon Bond and Futures Price

The volatility of the zero-coupon bond is given by $\sigma_P(t; T_P) \equiv \frac{\sigma_r}{\kappa_{rQ}} (1 - e^{-\kappa_{rQ}(T_P - t)})$, $\kappa_{rQ} \equiv \kappa_r + \beta_{rr}$.

The components of the volatility vector of the futures price are given by $\sigma_{Fr}(t; T_F) \equiv \frac{\sigma_r}{\kappa_{rQ}} (1 - e^{-\kappa_{rQ}(T_F - t)})$, $\sigma_{Fx}(t; T_F) \equiv \sigma_X$ and $\sigma_{r\delta}^{(i)}(t; T_F^{(i)}) \equiv -\frac{\sigma_\delta}{\kappa_{\delta Q}} (1 - e^{-\kappa_{\delta Q}(T_F - t)})$.

Appendix B: The Computation of the Certainty Equivalent

The deterministic functions needed to compute Equation (20.4b), for $n \in \{a, b\}$, are given by the following system of ordinary differential equations:

$$\begin{aligned} CE_{YY}^{(n)}(u) = & \frac{1}{\gamma} \lambda_Y^{(n)} \rho^{(n)-1} \lambda_Y^{(n)'} + [\mu_Y^{(n)} - g_Y \lambda_Y^{(n)} \sigma_Y^{(n)}] CE_{YY}^{(n)}(u) + CE_{YY}^{(n)}(u) [\mu_Y^{(n)} - g_Y \lambda_Y^{(n)} \sigma_Y^{(n)}] \\ & - g_Y CE_{YY}^{(n)}(u) \Sigma_Y^{(n)} CE_{YY}^{(n)}(u), \quad CE_{YY}^{(n)}(0) = 0^{(n,n)} \end{aligned} \quad (B.1)$$

$$\begin{aligned} CE_Y^{(n)}(u) = & e_2^{(n)} + \frac{1}{\gamma} \lambda_Y^{(n)} \rho^{(n)-1} \lambda_0^{(n)} + [\mu_Y^{(n)} - g_Y \lambda_Y^{(n)} \sigma_Y^{(n)}] CE_Y^{(n)}(u) + CE_{YY}^{(n)}(u) [\mu_0^{(n)} - g_Y \sigma_Y^{(n)} \lambda_0^{(n)}] \\ & - g_Y CE_{YY}^{(n)}(u) \Sigma_Y^{(n)} CE_Y^{(n)}(u), \quad CE_Y^{(n)}(u) = 0^{(n)} \end{aligned} \quad (B.2)$$

$$\begin{aligned} CE_0^{(n)}(u) = & \frac{1}{2\gamma} \lambda_0^{(n)'} \rho^{(n)-1} \lambda_0^{(n)} + [\mu_0^{(n)} - g_Y \sigma_Y^{(n)} \lambda_0^{(n)}]' CE_Y^{(n)}(u) - \frac{1}{2} g_Y CE_{YY}^{(n)}(u)' \Sigma_Y^{(n)} CE_Y^{(n)}(u) \\ & + \frac{1}{2} Tr[\Sigma_Y^{(n)} CE_{YY}^{(n)}(u)], \quad CE_0^{(n)}(0) = 0 \end{aligned} \quad (B.3)$$

where, for ease of exposition, we denoted by $g_Y \equiv 1 - 1/\gamma$, the weight invested in the portfolio $\pi_{CEht}^{(n)}$. $\Sigma_Y^{(n)} \equiv \sigma_Y^{(n)'} \rho^{(n)} \sigma_Y^{(n)}$ is the state variables matrix of the variance covariance, $E_t[Y_t^{(n)}] \equiv [\mu_0^{(n)} + \mu_Y^{(n)'} Y_t^{(n)}] dt$ is the state variables vector of expected return and $\lambda_t^{(n)} = \lambda_0^{(n)} + \lambda_Y^{(n)'} Y_t^{(n)}$ designates the vector of the market prices of risks. These quantities, necessary to compute the system of ordinary differential Equations (B.1)–(B.3), are given in the subappendixes, $B^{(n)}$; $n \in \{a, b\}$.

As far as the investment β is concerned, the system of ordinary differential equations needed to compute Equation (20.4b) is slightly different because of the market incompleteness that arises from the imperfect correlation between the spot commodity and its convenience yield:

$$\begin{aligned} CE_{YY}^{(\beta)}(u) = & \frac{1}{\gamma} \lambda_Y^{(\beta)} \rho^{(\beta)-1} \lambda_Y^{(\beta)'} + [\mu_Y^{(\beta)} - g_Y \lambda_Y^{(\beta)} \sigma_Y^{(\beta)}] CE_{YY}^{(\beta)}(u) \\ & + CE_{YY}^{(\beta)}(u) [\mu_Y^{(\beta)} - g_Y \lambda_Y^{(\beta)} \sigma_Y^{(\beta)}] - g_Y CE_{YY}^{(\beta)}(u) [\Sigma_Y^{(\beta)} - \gamma \cdot g_Y \hat{\Sigma}_{Y\beta}] CE_{YY}^{(\beta)}(u), \quad CE_{YY}^{(\beta)}(0) = 0^{(\beta,\beta)} \end{aligned} \quad (B.4)$$

$$\begin{aligned} CE_Y^{(\beta)}(u) = & e_2^{(\beta)} + \frac{1}{\gamma} \lambda_Y^{(\beta)} \rho^{(\beta)-1} \lambda_0^{(\beta)} + [\mu_Y^{(\beta)} - g_Y \lambda_Y^{(\beta)} \sigma_Y^{(\beta)}] CE_Y^{(\beta)}(u) \\ & + CE_{YY}^{(\beta)}(u) [\mu_0^{(\beta)} - g_Y \sigma_Y^{(\beta)} \lambda_0^{(\beta)}] - g_Y CE_{YY}^{(\beta)}(u) [\Sigma_Y^{(\beta)} - \gamma \cdot g_Y \hat{\Sigma}_{Y\beta}] CE_Y^{(\beta)}(u), \quad CE_Y^{(\beta)}(u) = 0^{(\beta)} \end{aligned} \quad (B.5)$$

$$CE_0^{(\beta)}(u) = \frac{1}{2\gamma} \lambda_0^{(\beta)'} \rho^{(\beta)-1} \lambda_0^{(\beta)} + [\mu_0^{(\beta)} - g_Y \sigma_Y^{(\beta)} \lambda_0^{(\beta)}]' CE_Y^{(\beta)}(u) - \frac{1}{2} g_Y CE_Y^{(\beta)}(u)' \Sigma_Y^{(\beta)} CE_Y^{(\beta)}(u) + \frac{1}{2} \text{Tr}[\Sigma_Y^{(\beta)} CE_{YY}^{(\beta)}(u)], \quad CE_0^{(\beta)}(0) = 0 \quad (\text{B.6})$$

The vector of the expected return of state variables $E_t[Y_t^{(\beta)}] \equiv [\mu_0^{(\beta)} + \mu_Y^{(\beta)'} Y_t^{(n)}] dt$ and their variance covariance matrix $\Sigma_Y^{(\beta)}$ are the same as for investment (b). However, because of market incompleteness, the matrix of correlation $\rho_Y^{(\beta)}$ and the market price of risk $\lambda_t^{(\beta)} = \lambda_0^{(\beta)} + \lambda_Y^{(\beta)'} Y_t^{(\beta)}$ differ. Moreover, a matrix that is linked to the risks of the convenience yield, $\hat{\Sigma}_{Y\beta}$, arises. These quantities, which are necessary to compute the system of ordinary differential Equations (B.4)–(B.6), are given in the subappendix B^(\beta).

B^(A) The Quantities Necessary to Compute Program (A)

For program (a), the vector of the state variables is given by the (log) equity price and the short rate, $Y_t^{(a)'} = [s_t \ r_t]$, which has dynamics that are governed by the correlated Brownian motion $dz_t^{(a)'} = [dz_{st} \ dz_{rt}]$,

with correlation matrix $\rho^{(a)} = \begin{bmatrix} 1 & \rho_{sr} \\ \rho_{sr} & 1 \end{bmatrix}$. $\sigma_Y^{(a)} = \begin{bmatrix} \sigma_s & 0 \\ 0 & \sigma_r \end{bmatrix}$ is the volatility matrix of the state variables.

$$\mu_0^{(a)} = \begin{bmatrix} -\sigma_s^2/2 + \beta_{s0} \\ \kappa_r \theta_r \end{bmatrix} \text{ and } \mu_Y^{(a)} = \begin{bmatrix} \beta_{ss} & 0 \\ 1 & -\kappa_r \end{bmatrix}.$$

The volatility of the traded assets, the equity and the zero-coupon bond of maturity T_p are given by

$$\sigma_A^{(a)}(t) = \begin{bmatrix} \sigma_s & 0 \\ 0 & -\sigma_p(t; T_p) \end{bmatrix}, \lambda_0^{(a)} = \begin{bmatrix} \lambda_{s0} \\ \lambda_{r0} \end{bmatrix}, \text{ and } \lambda_Y^{(a)} = \begin{bmatrix} \lambda_{ss} & 0 \\ 0 & \lambda_{rr} \end{bmatrix}.$$

By arbitrage, the expected returns of the traded assets are given by $\mu_{At}^{(a)} = r_t 1^{(a)} + \sigma_A^{(a)}(t)' \lambda_t^{(a)}$ with $1^{(a)'} = [1 \ 1]$. The proportions invested in the traded assets are $\pi_t^{(a)'} = [\pi_{st}^{(a)} \ \pi_{pt}^{(a)}]$, with $\pi_{st}^{(a)} \equiv \theta_{st}^{(a)} S_t / W_t^{(a)}$ and $\pi_{pt}^{(a)} \equiv \theta_{pt}^{(a)} P_t(T_p) / W_t^{(a)}$. $\theta_{st}^{(a)}$ and $\theta_{pt}^{(a)}$ are the number of units invested in the equity and bond markets, respectively. $e_2^{(a)}$ is equal to $e_2^{(a)'} = [0 \ 1]$. $0^{(a,a)}$ and $0^{(a)}$ designate the 2×2 null matrix and the two-dimensional null vector, respectively.

B^(B) The Quantities Necessary to Compute Program (B)

For program (b), the vector of state variables is given by the (log) equity price, the short rate, the (log) commodity spot price and its convenience yield $Y_t^{(b)'} = [s_t \ r_t \ x_t \ \delta_t]$, which has dynamics that are governed by the correlated Brownian motion $dz_t^{(b)'} = [dz_{st} \ dz_{rt} \ dz_{xt} \ dz_{\delta t}]$, with correlation matrix

$$\rho^{(b)} = \begin{bmatrix} 1 & \rho_{sr} & \rho_{sx} & \rho_{s\delta} \\ \rho_{sr} & 1 & \rho_{rx} & \rho_{r\delta} \\ \rho_{sx} & \rho_{rx} & 1 & \rho_{x\delta} \\ \rho_{s\delta} & \rho_{r\delta} & \rho_{x\delta} & 1 \end{bmatrix}, \sigma_Y^{(b)} = \begin{bmatrix} \sigma_s & 0 & 0 & 0 \\ 0 & \sigma_r & 0 & 0 \\ 0 & 0 & \sigma_x & 0 \\ 0 & 0 & 0 & \sigma_\delta \end{bmatrix} \text{ is the volatility matrix. } \mu_0^{(b)} = \begin{bmatrix} -\sigma_s^2/2 + \beta_{s0} \\ \kappa_r \theta_r \\ -\sigma_x^{(i)2}/2 + \beta_{x0}^{(i)} \\ \kappa_\delta^{(i)} \theta_\delta^{(i)} \end{bmatrix} \text{ and } \mu_Y^{(b)} = \begin{bmatrix} \beta_{ss} & 0 & 0 & 0 \\ 1 & -\kappa_r & 1 & 0 \\ 0 & 0 & \beta_{xx} & 0 \\ 0 & 0 & \beta_{x\delta} - 1 & -\kappa_\delta \end{bmatrix}.$$

The volatility of the traded assets, the equity, the zero-coupon bond of maturity T_p , the spot

commodity and the futures contract of maturity T_F are given $\sigma_A^{(b)}(t) = \begin{bmatrix} \sigma_s & 0 & 0 & 0 \\ 0 & -\sigma_p(t; T_p) & 0 & \sigma_{Fr}(t; T_F) \\ 0 & 0 & \sigma_x & \sigma_{Fx}(t; T_F) \\ 0 & 0 & 0 & \sigma_{F\delta}(t; T_F) \end{bmatrix}.$

$$\lambda_0^{(b)} = \begin{bmatrix} \lambda_{s0} \\ \lambda_{r0} \\ \lambda_{x0} \\ \lambda_{\delta 0} \end{bmatrix} \text{ and } \lambda_Y^{(b)} = \begin{bmatrix} \lambda_{ss} & 0 & 0 & 0 \\ 0 & \lambda_{rr} & 0 & 0 \\ 0 & 0 & \lambda_{xx} & 0 \\ 0 & 0 & \lambda_{x\delta} & \lambda_{\delta\delta} \end{bmatrix}. \text{ By arbitrage, the vector of the expected returns of the}$$

traded assets is given by $\mu_{At}^{(b)} = r_t 1^{(b)} + \sigma_A^{(b)}(t)' \lambda_t^{(b)}$ with $1^{(b)'} = [1 \ 1 \ 1 \ 0]$. The traded assets proportions are $\pi_t^{(b)'} = [\pi_{St}^{(b)} \ \pi_{Pt}^{(b)} \ \pi_{Xt}^{(b)} \ \pi_{\delta t}^{(b)}]$, with $\pi_{St}^{(b)} \equiv \theta_{St}^{(b)} S_t / W_t^{(b)}$, $\pi_{Pt}^{(b)} \equiv \theta_{Pt}^{(b)} P_t(T_p) / W_t^{(b)}$, $\pi_{Xt}^{(b)} \equiv \theta_{Xt}^{(b)} X_t / W_t^{(b)}$ and $\pi_{\delta t}^{(b)} \equiv \theta_{\delta t}^{(b)} F_t(T_f) / W_t^{(b)}$. $\theta_{St}^{(b)}, \theta_{Pt}^{(b)}, \theta_{Xt}^{(b)}$ and $\theta_{\delta t}^{(b)}$ are the number of units invested in the equity, bond, commodity and futures contracts, respectively. $e_2^{(b)}$ is equal to $e_2^{(b)'} = [0 \ 1 \ 0 \ 0]$. $0^{(b,b)}$ and $0^{(b)}$ designate the 4×4 null matrix and the four-dimensional null vector, respectively.

B^(β) The Quantities Necessary to Compute Program (β)

Program (β) differs from programs (a) and (b) because the traded assets cannot perfectly duplicate the risk faced by the investor, i.e. the risk of the state variables. Indeed, in this program, the investor faces the same economic environment as in program (b). As a consequence, the following relationships hold: $Y_t^{(\beta)} = Y_t^{(b)}$, $\sigma_Y^{(\beta)} = \sigma_Y^{(b)}$, $\Sigma_Y^{(\beta)} = \Sigma_Y^{(b)}$, $\mu_0^{(\beta)} = \mu_0^{(b)}$, $\mu_Y^{(\beta)} = \mu_Y^{(b)}$, $e_2^{(\beta)} = e_2^{(b)}$, $0^{(\beta,\beta)} = 0^{(b,b)}$ and $0^{(\beta)} = 0^{(b)}$. However, because of the market incompleteness, the correlation $\rho^{(\beta)}$ and the volatility matrixes $\sigma_Y^{(\beta)}$ differ and an additional term that relates to the risk of the convenience yield not hedged by the traded assets pertains:

$$\hat{\Sigma}_{Y\beta}. \text{ These three matrixes are given as follows: } \rho^{(\beta)} \equiv \begin{bmatrix} 1 & \rho_{sr} & \rho_{sX} \\ \rho_{sr} & 1 & \rho_{rX} \\ \rho_{sX} & \rho_{rX} & 1 \end{bmatrix}, \sigma_Y^{(\beta)} = \begin{bmatrix} \sigma_s & 0 & 0 & \zeta_{s\delta} \sigma_\delta \\ 0 & \sigma_r & 0 & \zeta_{r\delta} \sigma_\delta \\ 0 & 0 & \sigma_X & \zeta_{X\delta} \sigma_\delta \end{bmatrix}, \text{ with}$$

$$\begin{bmatrix} 1 & \rho_{sr} & \rho_{sX} \\ \rho_{sr} & 1 & \rho_{rX} \\ \rho_{sX} & \rho_{rX} & 1 \end{bmatrix} \begin{bmatrix} \zeta_{s\delta} \\ \zeta_{r\delta} \\ \zeta_{X\delta} \end{bmatrix} = \begin{bmatrix} \rho_{s\delta} \\ \rho_{r\delta} \\ \rho_{X\delta} \end{bmatrix}, \text{ and } \hat{\Sigma}_{Y\beta} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \zeta_{\delta\delta}^2 \sigma_\delta^2 \end{bmatrix} \text{ such that } 1 = \zeta' \rho^{(\beta)} \zeta + \zeta_{\delta\delta}^2 \text{ and } \zeta \equiv \begin{bmatrix} \zeta_{s\delta} \\ \zeta_{r\delta} \\ \zeta_{X\delta} \end{bmatrix}.$$

The volatility of the traded assets, the equity, the zero-coupon bond of maturity T_p and the spot commodity are given by $\sigma_A^{(\beta)}(t) = \begin{bmatrix} \sigma_s & 0 & 0 \\ 0 & -\sigma_p(t; T_p) & 0 \\ 0 & 0 & \sigma_X \end{bmatrix}$. The vector of Brownian motions that is associated

with this strategy is $dz_t^{(\beta)'} = [dz_{St} \ dz_{rt} \ dz_{Xt}]$. The parameters that are necessary to compute the market

$$\text{price of risk are } \lambda_0^{(\beta)} \equiv \begin{bmatrix} \lambda_{s0} \\ \lambda_{r0} \\ \lambda_{x0} \end{bmatrix} \text{ and } \lambda_Y^{(\beta)} \equiv \begin{bmatrix} \lambda_{ss} & 0 & 0 \\ 0 & \lambda_{rr} & 0 \\ 0 & 0 & \lambda_{xx} \\ 0 & 0 & \lambda_{x\delta} \end{bmatrix}. \text{ By arbitrage, the vector of expected returns for}$$

the traded assets is given by $\mu_{At}^{(\beta)} = r_t 1^{(\beta)} + \sigma_A^{(\beta)}(t)' \lambda_t^{(\beta)}$ with $1^{(\beta)'} = [1 \ 1 \ 1]$. The traded assets proportion is $\pi_t^{(\beta)'} = [\pi_{St}^{(\beta)} \ \pi_{Pt}^{(\beta)} \ \pi_{Xt}^{(\beta)}]$, with $\pi_{St}^{(\beta)} \equiv \theta_{St}^{(\beta)} S_t / W_t^{(\beta)}$, $\pi_{Pt}^{(\beta)} \equiv \theta_{Pt}^{(\beta)} P_t(T_p) / W_t^{(\beta)}$ and $\pi_{Xt}^{(\beta)} \equiv \theta_{Xt}^{(\beta)} X_t / W_t^{(\beta)}$. $\theta_{St}^{(\beta)}, \theta_{Pt}^{(\beta)}$ and $\theta_{Xt}^{(\beta)}$ are the number of units invested in the equity, bond and commodity markets, respectively.

Appendix C: (Log) Likelihood Functions

We display the (log) likelihood functions only in the case of a dynamic market price of risk. The case of the constant market price of risk is obtained by setting $\beta_{ss} = \beta_{rr} = \beta_{xx} = \beta_{x\delta} = \beta_{\delta\delta} = 0$ in the equations of Appendix C. In the following subappendixes, Δt stands for a time variation of one week and $N = 236$ stands for the number of dates of observations.

(Ca) Log Likelihood Function for the Equity Parameters

We apply it to lemma $s_t = \log(S_t)$ using Equation (20.1a) and discretize the obtained equation to get:*

$$s_{t+\Delta t} = [r_t + \beta_{s0}]\Delta t + [1 + \beta_{ss}\Delta t]s_t + \sigma_s dz_{st}, \quad (\text{Ca.1})$$

Equation (Ca.1) proves that the log equity price is conditionally Gaussian with variance σ_s^2 and with mean $E[s_{t+\Delta t}|s_t] = [r_t + \beta_{s0}]\Delta t + [1 + \beta_{ss}\Delta t]s_t$. By independence of conditional densities, we obtain the following pseudo likelihood function:

$$l(s_{\Delta t}, \dots, s_{N\Delta t}) \propto -(N-1) \log(\sigma_s \sqrt{\Delta t}) - \frac{1}{2\sigma_s^2 \Delta t} \sum_{i=1}^{i=N-1} (s_{(i+1)\Delta t} - E[s_{(i+1)\Delta t}|s_{i\Delta t}])^2, \quad (\text{Ca.2})$$

(Cb) Log Likelihood Function for the Interest Rate Parameters

Computation available from the author upon request shows that the (log) price of the bond can be couched as follows:

$$p_t(\tau_p) = \ln(P_t(T_p)) = B_{p0}(\tau_p) + B_{pr}(\tau_p)r_t, \tau_p = T_p - t \quad (\text{Cb.1})$$

$$\text{with } B_{pr}(\tau_p) = \frac{\sigma_r}{\kappa_{rQ}}(1 - e^{-\kappa_{rQ}\tau_p}), \text{ and } B_{p0}(\tau) = -\left[\theta_{rQ} - \frac{\sigma_r^2}{2(\kappa_{rQ})^2}\right]\left[\tau + B_{pr}(\tau_p)\right] - \frac{\sigma_r^2}{4\kappa_{rQ}}B_{pr}(\tau_p)^2.$$

The first principal component of the (log) term structure of bond prices is then given by:

$$p_t^{(1)} = B_{p0}^{(1)} + B_{pr}^{(1)}r_t \quad (\text{Cb.2})$$

where $B_{p0}^{(1)} = \sum_j \omega_{pj}^{(1)} B_{p0}(\tau_{pj})$, $B_{pr}^{(1)} = \sum_j \omega_{pj}^{(1)} B_{pr}(\tau_{pj})$ and the $\omega_{pj}^{(1)}$ are the loading factors of the first principal component.

We discretize the short rate in Equation (20.1b) as in subappendix (Ca) and use the change of variable given by Equation (Cb.2) to get the pseudo (log) likelihood function for the bond market:

$$l(p_{\Delta t}^{(1)}, \dots, p_{N\Delta t}^{(1)}) \propto -(N-1) \log(|B_{pr}^{(1)}| \sigma_r \sqrt{\Delta t}) - \frac{1}{2\sigma_r^2 \Delta t} \sum_{i=1}^{i=N-1} (r_{(i+1)\Delta t} - E[r_{(i+1)\Delta t} | r_{i\Delta t}])^2, \quad (\text{Cb.3})$$

where $E[r_{(i+1)\Delta t} | r_{i\Delta t}] = \kappa_r \theta_r \Delta t + [1 - \kappa_r \Delta t]r_{i\Delta t}$ and $r_{(i+1)\Delta t}$, $r_{i\Delta t}$ are given as functions of $p_{(i+1)\Delta t}^{(1)}$, $p_{i\Delta t}^{(1)}$, thanks to Equation (Cb.2).

(Cc) Log Likelihood Function for the Commodity Parameters

Because we estimate the parameters of the commodity markets independently of those of the bond markets, we choose to fit the (log) discounted futures prices: $f_i(\tau_F) = \log[P_t(T_F)F_i(T_F)]$; $\tau_F = T_F - t$. Because of the simplifying assumption of independence of the parameters, computation available from the author upon request shows that;

* Our discretization is not exact, but Schwartz (1997) shows that this simplification does not alter parameters estimation.

$$f_t(\tau_F) = x_t - D_{\kappa_{\delta Q}}(\tau_F)\delta_t + A(\tau_F) \quad (\text{Cc.1})$$

$$A(\tau_F) = \left(-\theta_{\delta Q} + \frac{\sigma_{\delta}^2}{2\kappa_{\delta Q}^2} - \frac{\sigma_{\delta}\sigma_X\rho_{X\delta}}{\kappa_{\delta Q}} \right) \tau_F + \frac{\sigma_{\delta}^2}{2\kappa_{\delta Q}^2} D_{2\kappa_{\delta Q}}(\tau_F) + \left(\theta_{\delta Q} - \frac{\sigma_{\delta}^2}{2\kappa_{\delta Q}^2} + \frac{\sigma_{\delta}\sigma_X\rho_{X\delta}}{\kappa_{\delta Q}} \right) D_{\kappa_{\delta Q}}(\tau_F)$$

and

$$D_{\kappa_{\delta Q}}(\tau_F) = \frac{1 - e^{-\kappa_{\delta Q}\tau_F}}{\kappa_{\delta Q}}.$$

We discretize Equations (20.1c)–(20.1d) as in the subappendix (Ca) to obtain:

$$\begin{aligned} \begin{bmatrix} x_{t+\Delta t} \\ \delta_{t+\Delta t} \end{bmatrix} &= \begin{bmatrix} r_t \Delta t + \left(\beta_{X0} - \frac{\sigma_X^2}{2} \right) \Delta t + (\beta_{X\delta} - 1) \Delta t \delta_t + (\beta_{Xx} \Delta t + 1) x_t \\ \sigma_X dz_{Xt} \kappa_{\delta} \theta_{\delta} \Delta t + \delta_t (1 - \kappa_{\delta} \Delta t) + \sigma_{\delta} dz_{\delta t} \end{bmatrix} \end{aligned} \quad (\text{Cc.1})$$

As a consequence, the (log) spot commodity price and its convenience yield are jointly conditionally

Gaussian with matrix of variance covariance $\Sigma = \begin{bmatrix} \sigma_{\delta}^2 & \sigma_X \sigma_{\delta} \rho_{X\delta} \\ \sigma_X \sigma_{\delta} \rho_{X\delta} & \sigma_X^2 \end{bmatrix} \Delta t$ and (conditional) mean:

$$E \begin{bmatrix} x_{t+\Delta t} \\ \delta_{t+\Delta t} \end{bmatrix} \begin{bmatrix} x_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} r_t \Delta t + \left(\beta_{X0} - \frac{\sigma_X^2}{2} \right) \Delta t \\ \kappa_{\delta} \theta_{\delta} \Delta t \end{bmatrix} + \begin{bmatrix} \beta_{Xx} \Delta t + 1 & (\beta_{X\delta} - 1) \Delta t \\ 0 & 1 - \kappa_{\delta} \Delta t \end{bmatrix} \begin{bmatrix} x_t \\ \delta_t \end{bmatrix} \quad (\text{Cc.2})$$

In the same vein as subappendix (Cb), we compute the first two principal components given by:

$$\begin{aligned} f_t^{(1)} &= \Omega^{(1)} x_t - D_{\kappa_{\delta Q}}^{(1)} \delta_t + A^{(1)} \\ f_t^{(2)} &= \Omega^{(2)} x_t - D_{\kappa_{\delta Q}}^{(2)} \delta_t + A^{(2)} \end{aligned} \quad (\text{Cc.3})$$

where $\Omega^{(k)} = \sum_j \omega_{fj}^{(k)}$, $D_{\kappa_{\delta Q}}^{(k)} = \sum_i \omega_{fj}^{(k)} D_{\kappa_{\delta Q}}(\tau_{Fj})$, $D_{\kappa_{\delta Q}}^{(k)} = \sum_i \omega_{fj}^{(k)} A(\tau_{Fj})$, $k = 1, 2$ and where $\omega_{fj}^{(1)}$, $\omega_{fj}^{(2)}$ are the loading factors of the first two principal components of the discounted futures prices, respectively.

We use a reasoning similar to that of subappendix (Cb) to get the pseudo (log) likelihood function:

$$\begin{aligned} l(f_{\Delta t}, \dots, f_{N\Delta t}) &\propto -(N-1) \log \left(\left| D_{\kappa}^{(1)} \Omega^{(2)} - \Omega^{(1)} D_{\kappa}^{(2)} \right| \sigma_X \sigma_{\delta} \sqrt{1 - \rho_{X\delta}^2} \Delta t \right) \\ &\quad - \frac{1}{2} \sum_{j=1}^{j=N-1} (Y_{(i+1)\Delta t} - E[Y_{(i+1)\Delta t} | Y_{i\Delta t}])' \Sigma^{-1} (Y_{(i+1)\Delta t} - E[Y_{(i+1)\Delta t} | Y_{i\Delta t}]) \end{aligned} \quad (\text{Cc.4})$$

where $f_{(i+1)\Delta t} = \begin{bmatrix} f_{(i+1)\Delta t}^{(1)} \\ f_{(i+1)\Delta t}^{(2)} \end{bmatrix}$, $f_{i\Delta t} = \begin{bmatrix} f_{i\Delta t}^{(1)} \\ f_{i\Delta t}^{(2)} \end{bmatrix}$, and $Y_{(i+1)\Delta t} = \begin{bmatrix} x_{(i+1)\Delta t} \\ \delta_{(i+1)\Delta t} \end{bmatrix}$, $Y_{i\Delta t} = \begin{bmatrix} x_{i\Delta t} \\ \delta_{i\Delta t} \end{bmatrix}$ are functions of $f_{(i+1)\Delta t}$, $f_{i\Delta t}$

thanks to Equation (Cc.3).

21

Long–Short Versus Long-Only Commodity Funds

21.1	Introduction	459
21.2	Sources of Alpha in Commodity Markets	459
21.3	Temporal Performance Measures	461
21.4	A Relative-Value Commodity Approach	462
21.5	Conclusions.....	466
	References.....	467

John M. Mulvey

21.1 Introduction

Commodity investments have gained considerable interest over the past decade. For traditionally diversified investors, an allocation to a fund that invests exclusively in commodity markets offers not only a hedge against inflation, but also effective diversification due to its low correlation with traditional asset classes. Over the long run, commodity investment funds show equity-like returns, but are accompanied by lower volatility and shortfall risk.

Fundamental flaws in traditional portfolio models became apparent during the severe 2008/09-crash period. Among other issues, many investors had assumed correlations in return among asset categories would approximate historical values. Under this assumption, the investor would be adequately and safely diversified to ‘protect’ capital. Unfortunately, most asset categories suffered together and even sophisticated investors, such as university endowment funds, lost substantial value—close to 30% for even the elite universities.

Absolute performance hedge funds failed to protect investor capital. In [Figure 21.1](#), we see that the correlation matrix among hedge fund categories became almost completely covered with ones. Facts of particular note are: (1) most hedge-fund categories lost money during this period; and (2) the Goldman Sachs Commodity Index (GSCI) suffered a precipitous drop of 51.4% over the specified period, including a maximum 75+% drawdown value. In [Figure 21.1](#), two categories—dedicated short bias and managed futures—stand out with very low correlation to the equity markets represented by the FTSE 500.

21.2 Sources of Alpha in Commodity Markets

The commodity segment of the managed futures domain can provide exceptional diversification from equities and fixed income. Commodity futures markets are among the oldest exchanges in the world, such as the Dojima rice futures market, which began in 1710 in Osaka, Japan. In recent years, investors have turned to owning commodities and other real assets to protect themselves against long-term risks.




FIGURE 21.1 Correlations among major hedge fund categories—Feb. 2008 to Feb. 2009 (heavy line = correlation > 0.5; light line = correlation between 0.2 and 0.5; no line = correlation < 0.2).

As the world population exceeds seven billion, the demand for basic commodities bumps against supply constraints for land, energy and agricultural products, possibly resulting in pricing disruptions. Even safe drinking water is becoming a scarce commodity in many parts of the world.

A related risk is inflation. Many countries are experiencing extraordinarily low nominal interest rates and massive deficit spending plans in order to overcome the fallout from the 2008/09 crash. The current level of negative real interest rates in a number of countries likely will contribute to future inflation.

Political risks, such as disruptions caused by oil embargos, wars and terrorist attacks, present another concern. Owning raw materials can be profitable during turbulent periods caused by political factors. Further, it is likely that equities will drop very quickly when a political crisis erupts. The 1973 oil embargo, for example, precipitated a substantial increase in energy prices, accompanied with higher inflation.

Last, there is a small, but still significant risk due to weather and catastrophic shocks such as crop freezes, hurricanes and tsunamis. Many commodity prices spike when these events occur.

Studies have shown the presence of trends and regime changes in commodity prices (Erb and Harvey 2006, Miffre and Rallis 2007, Shen *et al.* 2007). These patterns are due to the diffusion of information, inventory conditions, weather and political risks. Since most commodities are employed for consumption, either final or intermediate, consumers and producers render hedging decisions on an ongoing basis as a function of their core businesses. Many commodity-trading strategies are based on sustained price swings—either positive or negative—as a function of regime changes.

A second source of alpha involves the shape of the futures curve. In most commodities, the price of a futures contract is not determined by arbitrage. Rather, supply, demand and inventory considerations are paramount. Thus, for example, backwardation occurs when inventories are low and short-term

spikes in demand occur. Tactics based on the shape of the futures curve can lead to positive performance (Brennan *et al.* 1997, Gorton *et al.* 2008). A cause of the positive expected return for the traditional future curve strategy—selling in contango and buying in backwardation—goes back to the early work of Keynes around 1930. Therein, the general concept was that producers would be the predominant hedgers, thus resulting in backwardation as the normal course of events. In this situation, taking the other side of the bet—purchasing—would provide a positive expected performance, and empirical evidence supports this claim. Today, however, consumers such as the airlines and even pension plans who are attempting to hedge against inflation risks, have become the majority of hedgers, which results in a greater degree of contango. In this environment, there is a positive expected return for selling when the futures curve has severe contango (again empirical evidence supports this claim). Also see Bouchouev (2012) and Dempster *et al.* (2012) for further discussion.

There is controversy as to whether alpha is present in commodity strategies. Bhardwaj *et al.* 2008 indicates that commodity hedge funds rarely earn much more than the risk-free rate. This study was completed before the 2008/09 crash, however, wherein managed futures funds outperformed other categories by a wide margin (Figure 21.1). Lintner (1983) provides an early report showing the superior performance of selected commodity funds for improving performance relative to traditional assets.

21.3 Temporal Performance Measures

For long-term investors such as pension plans and university endowments, the most critical elements for their continued success are: (1) the growth of capital/wealth over time and (2) the investor's ability to support funding needs such as paying legal liabilities and other cash outflows. Traditional risk measures such as Sharpe ratios do not provide much information about the long-term financial health of the investor under a particular asset allocation mix. Instead, a multi-period portfolio model can be applied to determine more realistic future stochastic outcomes (e.g. Mulvey *et al.* (2008) and Dempster and Medova (2011)). In a multi-period context, it can be readily shown that long-term investors must prepare themselves to avoid a substantial loss of capital during periods of turbulence such as the 2008–09-crash period.

Given adequate historical evidence, we can construct temporal risk measures. One of the simplest is *maximum drawdown*. This value provides a limited starting point for analysing risks. Of course, max-drawdown, like most risk measures, is to some degree a backward-looking metric. Still, if we have observed a specified historical drawdown, we may well see a similar level of drawdown over future periods.

A related measure is the Ulcer index, which evaluates not only the amount of maximum drawdown, but also its duration. First, we define the historical wealth path: w_t : for $t = \{1, 2, \dots, T\}$, where the number of periods equals T . Next, the high water mark and drawdown are defined, respectively, as follows at each time period t : $h_t := \text{maximum } \{w_{t1}\}$ for $t1 = \{1, 2, \dots, t\}$ and $d_t = (h_t - w_t)$. Note that d_t is greater than or equal to zero.

Let u be the average drawdown over the entire historical period:

$$u = \sum_{t=1}^T d_t / T$$

Then the *Ulcer index* is the standard deviation of drawdown values:

$$Ulcer = \sum \left(\sum_{t=1}^T (d_t - u)^2 / T \right)^{1/2}$$

A related measure is the *Calmar index*—the average over fixed time periods, for instance 36 months, of return/the max-drawdown over this specified look-back period.

For highly structured decisions such as blackjack, the investor can measure the long-run growth of capital by reference to the Kelly investment strategy—by maximizing expected log (capital)—and its siblings. Unfortunately, the Kelly strategy can suffer relative long and sustained drawdown periods. Mulvey *et al.* (2011) combine the Kelly objective function with explicit drawdown constraints. A feature of multi-period models involves excess return over traditional static portfolio models. Herein, the level of excess return depends upon the volatility of the securities, diversification benefits and total transaction costs. In the case of commodity futures markets, volatility can be large, while transaction costs can be low—the ideal environment for achieving rebalancing gains. One of the simplest ways to achieve rebalancing gains is to apply a fixed-mix or even an equal-weighted portfolio of securities (see, for example, Dempster *et al.* (2011)). We will be evaluating the advantages of an equal-weighted portfolio of commodity tactics.

21.4 A Relative-Value Commodity Approach

Passive indexing strategies have become well established over the past 30 years. Index strategies are designed to match a well-defined market segment with a low cost (and possibly tax efficient) approach to investing. Passive indexing strategies can be pertinent for large institutional investors, due to their low cost and fees as well as their transparency and scalability.

Figures 21.2 and 21.3 show the wealth path for the S&P 500 equity index alongside two commodity indices—the Goldman Sachs Commodity Index (GSCI) and the Dow Jones/UBS index (DJUBS). These two commodity indices encompass long-only investments in the most actively traded commodities and serve as benchmarks for institutional investors.




FIGURE 21.2 Time series of two commodity indices and S&P 500 (1991–2012).




FIGURE 21.3 Time series of two commodity indices and S&P 500 (2000–2012).

The price patterns of the GSCI and DJUBS and equity markets such as the S&P 500 are roughly similar at first glance. The S&P 500 outperformed over the entire October 1991 to June 2012 period in terms of returns. Over the past 12 years, however, the results have flipped, with commodity indices beating the equity markets in terms of total return (Figure 21.3). It is clear that selecting a relevant time period is critical to relative comparisons.

In Figure 21.4, we see that the investible version of the GSCI (symbol GSG) achieved much lower performance than the spot index over the past 12 years. Underperformance is largely due to the presence of contango; commodities in contango produce lower returns due to the negative roll return. Several prominent commodities, including oil, have gone from backwardation to contango over the past decade as demand for these commodities has increased relative to their supply.

As mentioned, a critical issue for long-term investors involves drawdown risks. Unfortunately, both the GSCI and DJUBS experienced large drawdown values in 2008 equal to 70+%—at the exact time that equity markets crashed—and high Ulcer index values since drawdown continues.

To address drawdown and negative roll returns of long-only strategies, we propose a commodity portfolio via six established tactics: long and short momentum, long and short futures curve, trend following and breakout. Each of the first four tactics is based on a relative ranking of the commodities under study.

Briefly, we employ the six tactics to capture alpha embedded in commodity markets, while carefully balancing the long and short positions in the portfolio—in order to minimize drawdowns and produce positive returns with excellent diversification characteristics as compared with traditional assets (Mulvey 2012). Table 21.1 shows the performance of each of the individual six tactics over 1991 to 2012. Four of the tactics have roughly similar Sharpe ratios; however, adding the ‘inferior’ short tactics to a portfolio improves upon the overall performance over the full 1991 to 2012 time period.

A significant feature of futures (and forward) markets is the opportunity to be implemented as an overlay strategy without direct capital allocation. Herein, the investor employs their usual allocation to traditional assets such as stocks and bonds, e.g. 60/40. These assets provide margin requirements for the commodity futures contracts. Thus, the return of the commodity fund is additive for the investor. Especially important, therefore, is the diversification between traditional assets and the overlay strategies (discussed below).

First, we show the advantage of combining long and short directed funds. Figure 21.5 graphs the performance of the long and short momentum tactics. The long-momentum tactic outperforms short momentum over most of the entire span. However, during both crash periods—2001/02 and 2008/09—the short-momentum tactic outperformed its long-only counterpart by a wide margin. The two tactics combine 50/50 to provide a more stable return pattern—higher return per Ulcer Index. Remember that




FIGURE 21.4 Time series of GSCI index and investible ‘tracking’ fund (GSG)—source Bloomberg.

TABLE 21.1 Performance of Six Individual Tactics Over 1991–2012

Oct. 1991–Jun. 2012	Momentum Long	Momentum Short	Futures Curve Long	Futures Curve Short	Trend Following	Breakout
Geo. Return	8.76%	−0.25%	7.87%	4.64%	6.24%	6.51%
Volatility	15.21%	15.90%	15.42%	14.71%	9.21%	9.08%
Sharpe ratio**	0.38	−0.20	0.32	0.11	0.35	0.39
Drawdown	49.10%	52.16%	52.16%	48.64%	15.85%	15.89%
Ret/Drawdown	0.18	0.00	0.15	0.10	0.39	0.41
Ulcer index	15.19%	27.97%	15.18%	20.55%	5.80%	5.14%
UPI	0.58	−0.01	0.52	0.23	1.08	1.27

** 3% risk free rate.

Note: Data is back-tested only. All investors should be aware that future results may not be the same as historical performance.




FIGURE 21.5 Futures curve tactics (long and short) and market-neutral combination—1991 to 2012. (Results are from back test experiments. Readers should be aware of the limitations of results based solely on back tests performance.)

the cost of maintaining the low returning short-momentum tactic is modest due to the overlay structure; the short tactics provide low-cost ‘insurance’. The low Sharpe ratio of the short-momentum tactic understates the importance of this tactic for sound, long-term asset allocation.

A similar characteristic occurs with the long- and short-futures curve tactics. Again, the long-futures tactic has the better long-term performance as compared with the short-futures tactic, but does suffer from sharp drawdowns. Accordingly, the market-neutral 50/50 version has superior return/risk characteristics from the standpoint of return per drawdown or return per Ulcer Index. Table 21.2 lists the performance of a market neutral version (equal weighted) of the long and short versions of both the momentum and futures curve. The Ulcer index for the equal-weighted version is far superior to any of the four individual tactics—dropping below 4% versus 15–27% for the tactics. Thus, the market-neutral portfolio has a high Ulcer performance index (UPI)—return per Ulcer.


We next turn to evaluating a portfolio of the six mentioned tactics. Figure 21.6 and Table 21.3 depict the performance of an equal-weighted portfolio. In addition, we show that performance can be improved by applying regime detection for determining the tilting of long and short positions (Mulvey *et al.* 2011). We found that three regimes are most appropriate: (1) growth (66%), (2) transition (16%) and (3) contraction (18%). We tilt to the long (or short) side during regime 1 (or 3); the level of tilt is approximate 20%. The portfolio of six tactics improves performance over the previous four tactics since there are additional sources of uncertainties—leading to high rebalancing gains (e.g. Luenberger (1998), Dempster *et al.* (2003) and Mulvey *et al.* (2007)).

TABLE 21.2 Market-Neutral Combination of Momentum and Futures-Curve Tactics Versus SP 500

Oct. 1991– Jun. 2012	Momentum Long	Momentum Short	Futures Curve Long	Futures Curve Short	Market Neutral	Momentum Regime	S&P 500
Geo. Return	8.76%	–0.25%	7.87%	4.64%	6.24%	7.86%	5.81%
Volatility	15.21%	15.90%	15.42%	14.71%	6.15%	6.23%	18.78%
Sharpe ratio**	0.38	–0.20	0.32	0.11	0.53	0.78	0.15
Drawdown	49.10%	52.16%	52.16%	48.64%	9.99%	7.75%	56.78%
Ret/Drawdown	0.18	0.00	0.15	0.10	0.62	1.01	0.10
Ulcer Index	15.19%	27.97%	15.18%	20.55%	3.73%	2.87%	19.85%
UPI	0.58	–0.01	0.52	0.23	1.67	2.74	0.29

** 3% risk free rate.

Note: Data is back-tested only. All investors should be aware that future results may not be the same as historical performance.

**FIGURE 21.6** Wealth paths for long futures-curve tactic, equal-weighted combination and regime detection.**TABLE 21.3** Performance of a Portfolio of the Six Tactics

Oct. 1991–Jun. 2012	Equal Weight	Regime Detection
Geo. Return	6.36%	10.67%
Volatility	6.24%	8.60%
Sharpe ratio**	0.54	0.89
Drawdown	8.73%	10.36%
Ret/Drawdown	0.73	1.03
Ulcer index	3.35%	3.89%
UPI	1.90	2.74

** 3% risk free rate.

Note: Data is back-tested only. All investors should be aware that future results may not be the same as historical performance.

Last, we form a simplified combination of the commodity strategies—equal-weighted and regime-weighted—along with the S&P 500 equity index (Figure 21.7 and Table 21.4). Regime detection is based on a hidden Markov model, with several decades of out-of-sample performance (Fraser 2008, Mulvey *et al.* 2011, Prajogo 2011). The addition of the commodity portfolio as an overlay strategy greatly improves total return while modestly reducing drawdown values. The overlay strategies are additive to portfolio




FIGURE 21.7 Combining the commodity portfolios (equal-weighted and regime-weighted) as overlays with equities.

TABLE 21.4 Performance of Combined Equity and Commodity Fund

Oct. 1991–Jun. 2012	S&P 500	S&P 500 + Equal Weighted	S&P 500 + Regime
Geo. Return	5.81%	12.66%	17.16%
Volatility	18.78%	19.23%	20.37%
Sharpe ratio**	0.15	0.50	0.69
Drawdown	56.78%	45.06%	46.26%
Ret/Drawdown	0.10	0.28	0.37
Ulcer index	19.85%	14.31%	13.68%
UPI	0.29	0.67	1.04

** 3% risk free rate.

Note: Data is back-tested only. All investors should be aware that future results may not be the same as historical performance.

performance. Clearly, there are more desirable ways to combine commodities with traditional assets, but even this stylized version demonstrates the advantages of adding a long–short commodity fund to an equity portfolio.

21.5 Conclusions

This paper discusses the advantages of long–short commodity funds as meaningful diversifiers within a portfolio of traditional assets. The managed futures category of hedge funds, including commodity funds, performed particularly well during the 2008/09 crash periods. We have seen similar results in previous crash periods including the Asian currency crisis in 1997–98 and the Russian debt debacle and LTCM in 1998–99.

Positive performance during crashes can be attributed to: (1) the ready ability to go long or short; (2) deep liquidity and low transaction costs allowing for dynamic asset allocation and (3) the opportunity to take advantage of volatility via rebalancing gains and regime changes. Each element provides a small advantage. When combined, a portfolio of commodity tactics can substantially improve overall investment performance, especially when traditional assets are doing poorly.

The advantages of long-short commodity funds are described with attention to short positions during sharp downturns. We are currently working with FTSE to develop a dynamic index based on these principles (Mulvey 2012).

References

- Bhardwaj, G., Gorton, G.B. and Rouwenhorst, K.G., Fooling some of the people all of the time: the inefficient performance and persistence of commodity trading advisors. Working Paper 8(21), Yale ICF, 2008.
- Bodie, Z. and Rosansky, V.I., Risk and return in commodity futures. *Financ. Anal. J.*, 1980, **36**, 27–39.
- Bouchouev, I., The inconvenience yield or the theory of normal contango. *Quant. Finance*, 2012, this issue.
- Brennan, D., Williams, J. and Wright, B.D., Convenience yield without the convenience: a spatial-temporal interpretation of storage under backwardation. *Econ. J.*, 1997, **107**, 1009–1022.
- Dempster, M.A.H., Evstigneev, I.V. and Schenk-Hoppe, K.R., Exponential growth of fixed-mix strategies in stationary asset markets. *Finance Stochast.*, 2003, **7**, 263–276.
- Dempster, M.A.H., Evstigneev, I.V. and Schenk-Hoppe, K.R., Growing wealth with fixed mix strategies. In *The Kelly Capital Growth Investment Criterion*, edited by L. MacLean, E. Thorp and W. Ziemba, pp. 427–458, 2011 (World Scientific: Singapore).
- Dempster, M.A.H. and Medova, E., Asset-liability management for individual investors. *Br. Actuar. J.*, 2011, **16**(2), 405–464.
- Dempster, M.A.H., Medova, E. and Tang, K., Determinants of oil futures prices and convenience yields. *Quant. Finance*, 2012, this issue.
- Erb, C.B. and Harvey, C.R., The strategic and tactical value of commodity futures. *Financ. Anal. J.*, 2006, **62**(2), 69–97.
- Fraser, A.M., Hidden Markov models and dynamical systems. Society for Industrial and Applied Mathematics, 2008.
- Gorton, G.B., Hayashi, F. and Rouwenhorst, K.G., The fundamentals of commodity futures returns. Working Paper, Yale University, 2008.
- Gorton, G.B. and Rouwenhorst, K.G., Facts and fantasies about commodity futures. *Financ. Anal. J.*, 2006, **62**(2), 47–68.
- Greer, R.J., The nature of commodity index returns. *J. Altern. Invest.*, 2000, **3**, 45–52.
- Jeanneret, P., Monnin, P. and Scholz, S., Protection potential of commodity hedge funds. *J. Altern. Invest.*, 2011, **13**(3), 43–52.
- Lintner, J., The potential role of managed commodity-financial futures accounts (and/or funds) in portfolio of stocks and bonds. Working Paper, Harvard University, 1983.
- Luenberger, D., *Investment Science*, 1998 (Oxford University Press: Oxford).
- Miffre, J. and Rallis, G., Momentum in commodity futures markets. *J. Bank. Finance*, 2007, **31**(6), 1863–1886.
- Mulvey, J.M., The role of managed futures and commodity funds: protecting wealth during turbulent periods. *J. Indexes*, 2012, **15**(4), 38–45.
- Mulvey, J.M., Bilgili, M. and Vural, T., A dynamic portfolio of investment strategies: applying capital growth with drawdown penalties. In *The Kelly Capital Growth Investment Criterion*, edited by L. MacLean, E. Thorp and W. Ziemba, pp. 735–752, 2011 (World Scientific: Singapore).
- Mulvey, J.M., Simsek, K., Zhang, Z., Fabozzi, F. and Pauling, W., Assisting defined-benefit pension plans. *Oper. Res.*, 2008, **56**(5), 1066–1078.
- Mulvey, J.M., Ural, C. and Zhang, Z., Improving performance for long-term investors: wide diversification, leverage and overlay strategies. *Quant. Finance*, 2007, **7**(2), 1–13.
- Prajogo, A., Analyzing patterns in the equity market: ETF investor sentiment and corporate cash holdings. Ph.D. Dissertation, Princeton University, 2011.
- Shen, Q., Szakmary, A. and Sharma, S., An examination of momentum strategies in commodity futures markets. *J. Fut. Mkts*, 2007, **26**(3), 227–256.

22

Commodity Markets through the Business Cycle

Julien Chevallier
Mathieu Gatumel
Florian Ielpo

22.1	Introduction	469
22.2	The Reaction of Commodity Markets to Economic News.....	471
	Measuring the Impact of Price Discovery on Asset Prices • Database of News • An Example: S&P 500, US 10-Year and USD • Commodity by Commodity Analysis • Rolling Analysis	
22.3	Economic Regimes and Commodity Markets as an Asset Class	487
	Measuring the Business Cycle • To Which Business Cycle Are the Commodity Markets Related? • Commodity Performances Depending on the Nature of Each Economic Regime	
22.4	Conclusion	497
	References.....	497
	Appendix: Database of News	498

From 2008 to 2011, commodity markets experienced growing attention from the banking industry for various reasons: the summer 2008 oil price swing, the price surge in an ounce of gold, or sharp variations in agricultural prices. As a consequence, can we hypothesize the existence of a global connection between commodities and economic cycles? If these recent events suggest that commodity markets are strongly related to the business cycle, this evidence goes nevertheless against the widespread intuition that commodity markets are a strong source of diversification in a standard cash–bond–equity portfolio. Based on a data-set from 1990 to present, this paper investigates this issue by (i) looking at the reaction of commodity markets to economic news, and (ii) using a Markov regime-switching model to analyse economic regimes and commodity markets as an asset class.

Keywords: Commodities, Business cycle, News, Markov-switching

JEL Classification: G13, C63

22.1 Introduction

What is the relationship between the global macroeconomic environment and commodity markets (e.g. agricultural, metal and energy products)? From January 1987 to December 2010, the correlation between commodities and more traditional assets has been historically low: for example, the correlation between the WTI and the S&P 500 is equal to 3% on average. The correlation between gold and US government bonds can be rounded to 0%. Within commodities, the correlations are rather low as well in the long

run: for example, coffee and sugar are only related by a correlation coefficient of 8%. One of the largest correlations can be found between the WTI and gold—22%—which is much lower than the typical correlation found between the S&P500 and the Russell 3000 indices (equal to 85%).

In contrast with this view, it appears that since 2008, these correlation levels have been changing dramatically. In the long run, the correlation between the WTI and the Russell 3000 index was close to 5%. However, from 2008 to 2010, this correlation reached 28%—and the same comment arises for the S&P 500 index. Within commodities, a similar assessment can be made. Instead of the 8% coefficient found previously, the correlation between coffee and sugar is now closer to 35%.

What happened to the attractive ‘uncorrelation’ of commodity markets with other asset classes (Miffre and Rallis 2007, Geman and Ohana 2008, Daskalaki and Skiadopoulos 2011)?

One possible answer to this surge in correlation could be to assume some sort of stronger integration of these assets into global financial markets: the growing inclusion of commodity futures in strategies built by hedge funds, or in the balanced mandates managed by asset managers, may impact the behaviour of these futures with respect to other markets. How can we measure this integration phenomenon? How integrated now are commodity markets with equity and bond markets?

This paper aims at gathering different pieces of evidence to contribute to a better understanding of these key questions. Indeed, previous literature has already documented that some cross-market linkages do exist between commodity markets and the global macroeconomic environment (Caballero *et al.* 2008, Chng 2009, Tang and Xiong 2010, Chan *et al.* 2011, Dionne *et al.* 2011, Chevallier 2012). Besides, commodity price modelling has emerged as a vivid area of research in the financial literature (Miltersen 2003, Tang 2012).

First, if an increasing integration between commodity and more traditional asset markets has been taking place over the past 10 years, this must be measurable by the reaction of these markets to economic news. Thus, we develop a commodity by commodity analysis of the impact of news. This analysis has been performed recently by Elder *et al.* (2012) for metal futures only. Our results reveal an interesting pattern: the response of commodity prices to economic surprises is strong during global downturns, and weak during expansion periods. In addition, by slicing the information available into two-year subsets, the results show a very limited support to the ‘financial integration’ theory. Our results suggest that commodity markets have been over-reacting to economic news during the 2008–2009 crisis, similarly to what happened in 2001. During 2009–2010, our estimations reveal indeed a decreased sensitivity of commodity markets to the economic newsflow.

Second, instead of evaluating how commodity markets are affected by economic surprises, we seek to discuss the relationship between the business cycle in various economic zones and the performance of the different commodity sectors. Namely, we investigate whether commodity markets are related to a given economic zone or to the worldwide business cycle—mainly through the global demand for raw materials and energy. To do so, we consider how commodity prices evolve through the business cycle based on the class of Markov regime-switching models. This approach has been applied successfully to commodity markets in previous literature (Alizadeh *et al.* 2008, Andersen 2010). We apply this methodology for the USA, the Eurozone and China. Our results cast some light on the strong relationship that appears between commodity markets and the underlying business cycle. More especially, we are able to detect an increased sensitivity to economic activity in China.

Moreover, we provide a more detailed analysis of the US business cycle by decomposing its phases based on the joint behaviour of production growth, inflation, unemployment and retail sales behaviour. The US cycle is broken down into five different types of regimes: strong expansion, medium expansion, stalling, slowdown and strong crisis. For each of these regimes, we characterize the performance of commodities by assessing which phase coincides with rising or falling commodity prices. Besides, we identify during which periods commodity markets are most likely to grow or fall.

The remainder of the paper is structured as follows. [Section 22.2](#) details the incorporation of news in commodity markets. [Section 22.3](#) relates commodity markets to the business cycles. [Section 22.4](#) briefly concludes.

22.2 The Reaction of Commodity Markets to Economic News

The investigation of the reaction of commodity markets to economic news is based on the intuition that liquid financial markets integrate efficiently new information.

22.2.1 Measuring the Impact of Price Discovery on Asset Prices

One way to determine what is priced on a given market is to measure the impact of news on asset prices. Note that $R_{i,t}$ is the actual value for the release figure i at time t and $F_{i,t}$ is the market consensus for this figure.* The ‘surprise’ component in this economic news $S_{i,t}$ is given by:

$$S_{i,t} = R_{i,t} - F_{i,t}. \quad (22.1)$$

Given that various news may have different scales, it is very common to scale $S_{i,t}$ by its full sample standard deviation. By scaling the surprises by $\sigma_i = \sqrt{V(S_{i,t})}$, we make the various announcements comparable and it will thus be possible to rank the main movers on commodity markets. Denote $s_{i,t}$ the scaled variable:

$$s_{i,t} = \frac{R_{i,t} - F_{i,t}}{\sigma_i}. \quad (22.2)$$

For a given financial asset, the return on investment on day t is denoted r_t . This sampling frequency allows us to focus on the impact of surprises on the fundamental value of the asset studied.† Formally, the news analysis is performed by running the following sets of regressions:

$$r_t = \phi_0 + \phi_1 r_{t-1} + \sum_{i=1}^n \mathbb{I}_i \beta_i \times s_{i,t} + \sigma_t \varepsilon_t \quad (22.3)$$

$$\log \sigma_t^2 = \omega + \alpha \varepsilon_{t-1} + \theta |\varepsilon_{t-1}| + \log \sigma_{t-1}^2. \quad (22.4)$$

with $\mathbb{I}_i$ a dummy variable equal to 1 whenever the news i is released at time t , and 0 otherwise n is the total number of news. The analysis focuses mainly on the sensitivity of r_t to $s_{i,t}$, that is to say β_i . Note that:

- Whenever β_i is statistically different from 0, it seems reasonable to assume that this asset’s price incorporates the economic information associated with the surprise component $s_{i,t}$.
- The sign of β_i is essential to the market interpretation of the data.
- Equations (22.3) and (22.4) incorporate two types of robustness checks: an autoregressive component through $\phi_0 + \phi_1 r_{t-1}$, and a time-varying volatility component through the dynamics of the exponential GARCH (EGARCH) model.

Overall, this methodology corresponds to standard practice in order to gauge the impact of news on financial returns. We propose to apply it to the analysis of commodity markets.

* The series $F_{i,t}$ is taken from the Bloomberg survey data-set. It represents the median value of a survey of economists. Given that Bloomberg is one of the most widespread information services used by practitioners and market participants, it seems reasonable to use this data-set.

† Although an approach based on a higher frequency can also be interesting to commodity investors, a price variation that would not last at least one day can hardly be regarded as a fundamental change in the value of the investigated asset.

22.2.2 Database of News

In this section, we present the Bloomberg data-set of 16 series of surprises stemming from three economic zones: the USA, the EMU and China. For the news associated to Europe and the USA, the data-set starts in 1999 and ends in 2011. For the news associated to China, the data-set starts in 2007 and ends in 2011.

The news cover different aspects of the business cycle: real activity with macroeconomic data (such as industrial production) or surveys (such as the ISM or the German IFO), inflation dynamics with the consumer price index (CPI), monetary variables with the Fed Target Rate and the ECB's minimum refinancing rate.*

A detailed list of these news is provided in Table 22.1, along with the correlation between the US GDP and each of these news. The interested reader can find additional background information in the Appendix.

Except for the weekly Jobless Claims—that is the weekly number of new jobless applications in the USA—and the ECB decision rate, the majority of economic news is found to have a positive correlation with the US business cycle (as measured by GDP growth).

The weakest positive correlation is obtained for the German ZEW: this survey gathers the opinion of market participants regarding the current state of the German economy. Another German survey—the IFO survey—is computed in a very different way, by collecting the opinion of purchasing managers. In this respect, the IFO is closer in its spirit to the famous US ISM survey, that is why we find a strong correlation with the US GDP in the latter case.

Before turning to the analysis of commodities, we illustrate this methodology by showing the estimation results obtained when using traditional assets.

22.2.3 An Example: S&P 500, US 10-Year and USD

Before dealing with the impact of news on commodity markets, we start first with a review of the impact of news on standard assets: the S&P 500 index, the US 10-year bond rate and the trade weighted value

TABLE 22.1 Database of Economic News From Bloomberg

	Type of data	Geographical zone	Correlation with US GDP (%)
Non-Farm Payroll	Employment	USA	77
ISM	Economic survey	USA	84
Jobless Claims	Employment	USA	−68
US CPI MoM	Inflation	USA	22
US Retail Sales	Economic activity	USA	27
Fed Target Rate	Central bank rate	USA	3
US GDP	Economic activity	USA	100
ZEW Eco. Sent.	Economic survey	Germany	20
IFO Expectations	Economic survey	Germany	75
EMU CPI	Inflation	EMU	26
EMU GDP	Economic activity	EMU	80
FR Bus. Conf.	Economic survey	France	70
ECB Ref. Rate	Central bank rate	EMU	−13
China CPI YoY	Inflation	China	45
China Ind. Prod.	Economic activity	China	73
China PMI	Economic survey	China	41

* Note that the Chinese decision rate is not modified through a prescheduled decision meeting, as in the case of the Fed and the ECB. For this reason—and due also to the lack of consistency that it would introduce in the analysis—we have discarded the announcements by the Public Bank of China (PBOC).

of the US dollar. This preliminary step can be seen as a benchmark of comparison with commodities. Table 22.2 contains the estimation analysis obtained with these three assets.

Out of the seven US economic indicators, the 10-year US bond rate responds to six of them, i.e. to all figures except the GDP itself. The US dollar index also responds positively to the ISM, the Non-Farm Payrolls and the Fed decision rate. The S&P 500 responds positively to the retail sales, and negatively to the Fed action. When strictly focusing on the US economy, the equity index is clearly less sensitive to economic indicators than the 10-year bond rate and the US dollar. Besides, equity markets are also prone to react to economic news coming from other countries: the S&P 500 reacts positively to the ECB decision rate, and negatively to European news such as the ZEW survey (unlike the US dollar which reacts negatively to good news coming from the ZEW survey). Finally, the US dollar is found to react negatively to a positive surprise in the French business confidence index, and positively to an unexpected increase in the CPI.

Next, we investigate in more detail the behaviour of commodity markets.

22.2.4 Commodity by Commodity Analysis

In this section, we report the results obtained when estimating Equations (22.3)–(22.4) commodity by commodity.

22.2.4.1 Database for Commodity Prices

Table 22.3 gives the characteristics of each commodity futures price series included in the news analysis. Closest-to-maturity futures employed, in the sense that rollover returns are taken into account from the expiring contract to the one with the second-to-maturity contract to avoid the so-called ‘Samuelson effect’ (with a rollover date 15 days before maturity). The data comes from Bloomberg with a monthly frequency.*

TABLE 22.2 Impact of News on the S&P 500 Index, the US 10-Year Rate and the Trade Weighted Value of the US Dollar

	S&P 500	10Y US	US Dollar
Intercept	−7.15** (−4.68)	−6.35** (−4.21)	−1.25 (−0.82)
AR	0.02 (1.23)	−0.02 (−1.05)	0 (0.07)
Non-Farm Payroll	−0.09 (−0.99)	0.75** (6.31)	0.19** (5.07)
ISM	0.1 (1.01)	0.55** (4.62)	0.13** (3.29)
Jobless Claims	−0.03 (−0.34)	−0.33** (−2.82)	−0.04 (−1.04)
US CPI MoM	−0.11 (−1.19)	0.4** (3.3)	0.04 (0.91)
US Retail Sales	0.29** (2.59)	0.7** (4.92)	0.07 (1.42)
Fed Target Rate	−0.25* (−1.86)	0.4** (2.29)	0.19** (3.4)
US GDP	0.06 (0.33)	−0.21 (−0.96)	−0.02 (−0.26)
ZEW Economic Sentiment	−0.29** (−2.46)	−0.04 (−0.24)	0.08* (1.76)
IFO Expectations	−0.03 (−0.22)	0.16 (0.96)	−0.04 (−0.85)
EMU CPI Flash Estimate	−0.11 (−0.9)	0 (0.01)	0.03 (0.73)
EMU GDP	0.13 (0.71)	−0.08 (−0.32)	0.03 (0.41)
France Business Confidence	0.05 (0.45)	0.21 (1.47)	−0.09* (−1.84)
ECB Refinancing Rate	0.37** (3.32)	−0.02 (−0.16)	0.03 (0.63)
China CPI YoY	0.02 (0.14)	−0.05 (−0.33)	0.08* (1.7)
China Industrial Production	−0.01 (−0.11)	0.03 (0.19)	−0.02 (−0.51)
China PMI	0.08 (0.51)	0.05 (0.27)	0.02 (0.4)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

* This comment applies in the remainder of the paper.

TABLE 22.3 Database of Commodity Futures Used for the News Analysis

Commodity	Exchange	Unit
Gold	COMEX Gold Futures price	US Dollar per troy ounce
Silver	COMEX Silver Futures price	US Dollar per troy ounce
Platinum	COMEX Platinum Futures price	US Dollar per troy ounce
Aluminium	LME Aluminium Futures price	US Dollar per tonne
Copper	COMEX Copper Futures price	US Dollar per pound
Nickel	LME Nickel Futures price	US Dollar per tonne
Zinc	LME Zinc Futures price	US Dollar per tonne
Lead	LME Lead Futures price	US Dollar per tonne
WTI	NYMEX WTI Light Sweet Crude Oil Futures price	US Dollar per barrel
Brent	ICE Brent Crude Futures price	US Dollar per barrel
Gasoil	RBOB Gasoline Futures price	US Dollar per gallon
Natural Gas	Henry Hub Natural Gas Futures price	US Dollar per mmBtu
Heating Oil	NYMEX Heating Oil Futures price	US Dollar per gallon
Corn	CBOT Corn Futures price	US Dollar per bushel
Wheat	CBOT Wheat Futures price	US Dollar per bushel
Coffee	ICE Coffee C Futures price	US Dollar per pound
Sugar	NYMEX Sugar No.11 Futures price	US Dollar per pound
Cocoa	ICE Cocoa Futures price	US Dollar per tonne
Cotton	ICE Cotton No.2 Futures price	US Dollar per pound
Soybean	CBOT Soybean Futures price	US Dollar per short ton
Rice	CME Rough Rice Futures price	US Dollar per hundred weight

Source: Bloomberg.

22.2.4.2 Precious Metals

Tables 22.4 and 22.5 present the results obtained with a subset of precious metals, i.e. gold, silver and platinum. When considering the full sample, we obtain the results reproduced in Table 22.4. When considering the potential reaction of precious metals depending on the phase of business cycle—i.e. ‘expansion’ or ‘recession’—as indicated by the NBER Business Cycle Dating Committee,* we obtain the results reproduced in Table 22.5. In each table, the Goldman Sachs Commodity Index (GSCI) for precious metals is used as a benchmark.†

The results may be summarized as follows:‡

1. *Gold*: From a long-term perspective, gold reacts negatively to positive surprises in Non-Farm Payrolls and the Chinese PMI. In this respect, sustained and accelerating growth regimes are thus likely to hinder the rise of gold prices. We do not find evidence of a reaction of gold prices to surprises concerning inflation news. When breaking down the phases of the business cycle (Table 22.5), we obtain a different picture. During recessions, gold reacts negatively to the US ISM, the EMU CPI and the Chinese PMI. Two out of the three figures having a statistical impact on gold are economic surveys. These surveys are to be monitored cautiously to anticipate periods of rising gold prices. During expansion periods, gold displays a negative sign with the Non-Farm Payrolls, but a positive sign with the US CPI and the German IFO. Hence, during

* The NBER methodology of dating business cycle is authoritative and well known, and widely used by financial practitioners as well.

† The GSCI is a broadly diversified index designed for investors seeking to store value in commodity markets.

‡ To conserve space, we skip the results of the EGARCH model, where all parameters are highly significant. They can be obtained upon request to the authors.

TABLE 22.4 Estimation of the Impact of Economic News Over Selected Precious Metals

	Gold	Silver	Platinum	GSCI Prec. Metals
Intercept	-1.89 (-1.25)	-1.73 (-1.15)	4.11** (2.73)	1.52 (0.99)
AR	0.03** (2.16)	0.04 (1.59)	0.03 (1.2)	0.04** (2.19)
Non-Farm Payroll	-0.17** (-2.16)	-0.18 (-1.32)	-0.18* (-1.7)	-0.16* (-1.92)
ISM	-0.03 (-0.37)	0.05 (0.39)	0.07 (0.7)	-0.03 (-0.32)
Jobless Claims	-0.05 (-0.59)	-0.27* (-1.91)	-0.1 (-0.96)	-0.05 (-0.57)
US CPI MoM	0.06 (0.71)	0.11 (0.75)	0.11 (1.06)	0.03 (0.38)
US Retail Sales	-0.03 (-0.31)	-0.04 (-0.23)	-0.02 (-0.13)	-0.1 (-1.01)
Fed Target Rate	-0.04 (-0.33)	-0.23 (-1.12)	-0.09 (-0.58)	-0.02 (-0.19)
US GDP	-0.01 (-0.04)	0.08 (0.29)	0.38** (1.98)	-0.03 (-0.21)
ZEW Eco. Sent.	0.07 (0.76)	0.03 (0.16)	-0.21 (-1.6)	0.09 (0.86)
IFO Expectations	0.13 (1.22)	0.18 (0.92)	0.2 (1.41)	0.08 (0.72)
EMU CPI	-0.05 (-0.46)	-0.2 (-1.14)	-0.18 (-1.37)	-0.09 (-0.85)
EMU GDP	-0.08 (-0.49)	-0.11 (-0.37)	-0.04 (-0.19)	-0.08 (-0.49)
FR Bus. Conf.	0.12 (1.32)	0.02 (0.15)	0.1 (0.8)	0.13 (1.25)
ECB Ref. Rate	-0.06 (-0.67)	0.14 (0.87)	0.12 (0.99)	-0.07 (-0.7)
China CPI YoY	-0.03 (-0.33)	0.04 (0.21)	-0.06 (-0.51)	-0.1 (-0.98)
China Ind. Prod.	0 (0.03)	0.01 (0.08)	0.26** (2.18)	0.05 (0.54)
China PMI	-0.43** (-3.58)	-0.48** (-2.23)	-0.23 (-1.45)	-0.5** (-3.92)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

expansion, positive surprises concerning inflation—i.e. when the inflation index rises more than expected—should trigger a surge in gold prices. Overall, we identify that gold has a very different reaction to economic news depending on the underlying economic regime.

2. *Silver*: From a long-term perspective, silver reacts negatively to positive surprises in the Jobless Claims and the Chinese PMI. The picture is more complex than for gold: the negative sign with the Chinese PMI underlines the ‘safe haven’ role of silver during recessions, whereas its negative reaction to an unexpected increase in the number of unemployed people in the USA tells us a different story. This effect comes from the key fact that many precious metals are used both as safe havens (Baur and McDermott 2010, Beaudry *et al.* 2011) and for industrial purposes. During recession periods, silver reacts positively to negative surprises in the Chinese PMI, and negatively to positive surprises in the French business confidence indicator. During expansion periods, silver reacts positively to positive surprises in the US CPI, the Chinese PMI and the IFO Expectations. It has a weaker sensitivity to the Jobless Claims.
3. *Platinum*: During the full sample (Table 22.4), platinum shows a negative sensitivity to Non-Farm Payroll, the German ZEW and the EMU CPI, and a positive reaction to the US GDP and the Chinese industrial production. The most striking fact for platinum is that it features a stronger relationship to economic news than gold and silver: while gold only reacts to two news, platinum reacts to five of the news considered. We also clearly see the mix between the safe haven characteristic of the average precious metal, and the apparent use of such metal for industrial purposes. When considering only recession periods, the strongest sensitivity of platinum is obtained for the US GDP: an unexpected decrease in the US GDP would lead to a strong decline in the prices of platinum. Two minor effects come from the sign of the German ZEW and the EMU CPI, in a contracyclical way. Expansion periods are characterized by a rise in the platinum price, whenever the US CPI and the Chinese industrial production surprises are on the upside. Similarly to the GSCI benchmark for precious metals, platinum can be used for industrial purposes. Hence, it reacts positively to positive news about economic growth, and it is traditionally used as a hedge against inflation.

TABLE 22.5 Estimation of the Impact of Economic News Over Selected Precious Metals Depending on the NBER Phase of the Business Cycle

		Gold	Silver	Platinum	GSCI Prec. Metals
REC	Intercept	-2.15 (-1.41)	-1.66 (-1.08)	4.38** (2.88)	1.29 (0.84)
	AR	0.03** (2.2)	0.05* (1.64)	0.03 (1.23)	0.04** (2.2)
	Non-Farm Payroll	0.03 (0.17)	-0.32 (-0.86)	-0.42 (-1.51)	0.07 (0.33)
	ISM	-0.36** (-2.04)	-0.49 (-1.54)	-0.34 (-1.41)	-0.41** (-2.19)
	Jobless Claims	-0.06 (-0.45)	-0.24 (-1.02)	-0.18 (-1.01)	-0.01 (-0.04)
	US CPI MoM	-0.19 (-1.29)	-0.36 (-1.37)	-0.24 (-1.23)	-0.22 (-1.41)
	US Retail Sales	-0.02 (-0.15)	0.02 (0.08)	0.17 (1.03)	-0.1 (-0.78)
	Fed Target Rate	0.01 (0.05)	-0.17 (-0.66)	-0.12 (-0.6)	0.02 (0.1)
	US GDP	0.1 (0.24)	0.22 (0.3)	1.98** (3.54)	-0.05 (-0.11)
	ZEW Eco. Sent.	0.29 (1.44)	-0.26 (-0.73)	-0.92** (-3.41)	0.12 (0.58)
	IFO Expectations	-0.1 (-0.55)	-0.41 (-1.21)	0.15 (0.61)	-0.36* (-1.81)
	EMU CPI	-0.39** (-1.99)	-0.49 (-1.41)	-0.71** (-2.75)	-0.36* (-1.78)
	EMU GDP	-0.62 (-1.18)	-0.8 (-0.85)	-0.73 (-1.05)	-0.48 (-0.88)
	FR Bus. Conf.	0.6** (3.72)	0.74** (2.58)	0.3 (1.38)	0.53** (3.11)
	ECB Ref. Rate	-0.06 (-0.55)	0.09 (0.44)	0.15 (0.99)	-0.05 (-0.4)
	China CPI YoY	0.36 (1.14)	0.45 (0.79)	-0.01 (-0.03)	-0.28 (-0.84)
	China Ind. Prod.	-0.11 (-0.51)	0.25 (0.62)	0.09 (0.29)	0.04 (0.19)
	China PMI	-0.51** (-3.88)	-0.66** (-2.82)	-0.25 (-1.44)	-0.59** (-4.26)
EXP	Non-Farm Payroll	-0.21** (-2.49)	-0.17 (-1.15)	-0.14 (-1.31)	-0.2** (-2.29)
	ISM	0.06 (0.67)	0.18 (1.19)	0.17 (1.51)	0.07 (0.77)
	Jobless Claims	-0.02 (-0.19)	-0.24 (-1.41)	-0.04 (-0.32)	-0.05 (-0.49)
	US CPI MoM	0.16* (1.74)	0.31* (1.83)	0.26** (2.07)	0.13 (1.37)
	US Retail Sales	-0.04 (-0.31)	-0.1 (-0.4)	-0.25 (-1.3)	-0.09 (-0.61)
	Fed Target Rate	-0.03 (-0.17)	-0.17 (-0.52)	0.05 (0.23)	0 (0.01)
	US GDP	0 (0)	0.08 (0.28)	0.18 (0.89)	-0.02 (-0.12)
	ZEW Eco. Sent.	0.01 (0.09)	0.11 (0.57)	0 (-0.02)	0.07 (0.6)
	IFO Expectations	0.22* (1.68)	0.42* (1.81)	0.23 (1.3)	0.28** (2.01)
	EMU CPI	0.1 (0.88)	-0.05 (-0.25)	0.03 (0.17)	0.03 (0.27)
	EMU GDP	-0.02 (-0.14)	-0.04 (-0.12)	0.03 (0.13)	-0.04 (-0.22)
	FR Bus. Conf.	-0.12 (-1.03)	-0.34 (-1.6)	0 (0.01)	-0.07 (-0.61)
	ECB Ref. Rate	-0.1 (-0.66)	0.16 (0.59)	0.06 (0.32)	-0.14 (-0.91)
	China CPI YoY	-0.07 (-0.75)	-0.02 (-0.12)	-0.1 (-0.74)	-0.09 (-0.83)
	China Ind. Prod.	0.02 (0.22)	-0.04 (-0.2)	0.29** (2.24)	0.05 (0.48)
	China PMI	0.38 (1.11)	1.03* (1.66)	0.4 (0.87)	0.34 (0.92)

Notes: REC stands for 'Recession' and EXP for 'Expansion' according to the NBER Business Cycle Dating Committee. T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

22.2.4.3 Industrial Metals

In the case of industrial metals, we use the following subset of markets: Aluminium, Copper, Nickel, Zinc and Lead. As for precious metals, we report the estimation results obtained during the full sample (Table 22.6), and during expansion vs. recession periods (Table 22.7).

Our results can be summarized as follows:

1. *Aluminium*: Over the 1999–2011 period, the price of aluminium has been rising with Non-Farm Payrolls, the US ISM and ECB rates. On the contrary, it reacted negatively to increases in the number of unemployed people in the US, as measured by the Jobless Claims. Hence, in the long

TABLE 22.6 Estimation of the Impact of Economic News Over Selected Industrial Metals

	Aluminium	Copper	Nickel	Zinc	Lead	GSCI Ind. Metals
Intercept	-3.04** (-1.98)	-4.7** (-3.07)	0.71 (0.46)	-2.61* (-1.71)	5.39** (3.52)	-4.53** (-2.95)
AR	0 (0.21)	0.03 (1.13)	0.02 (0.64)	0.01 (0.5)	0.03 (0.86)	0.02 (1.07)
Non-Farm Payroll	0.22** (2.32)	0.15 (1.19)	0.16 (0.97)	0.2 (1.52)	0.09 (0.59)	0.19* (1.8)
ISM	0.16* (1.65)	0.37** (2.96)	0.24 (1.39)	0.18 (1.37)	0.21 (1.48)	0.21** (1.99)
Jobless Claims	-0.18* (-1.89)	-0.28** (-2.26)	-0.32* (-1.92)	-0.21 (-1.59)	-0.2 (-1.38)	-0.24** (-2.33)
US CPI MoM	-0.11 (-1.17)	0.04 (0.29)	-0.28 (-1.63)	-0.12 (-0.86)	0.03 (0.2)	-0.06 (-0.61)
US Retail Sales	0.16 (1.4)	-0.04 (-0.29)	-0.53** (-2.63)	0.02 (0.12)	-0.22 (-1.24)	0 (-0.03)
Fed Target Rate	0.05 (0.33)	0.03 (0.16)	0.05 (0.2)	0.05 (0.26)	0.27 (1.28)	0.08 (0.51)
US GDP	0.12 (0.69)	0.37 (1.61)	0.02 (0.06)	0.5** (2.01)	-0.05 (-0.17)	0.24 (1.25)
ZEW Eco. Sent.	-0.12 (-0.99)	-0.27* (-1.73)	-0.21 (-1)	0.16 (0.97)	-0.24 (-1.3)	-0.2 (-1.56)
IFO Expectations	-0.18 (-1.36)	-0.26 (-1.54)	-0.27 (-1.14)	0 (0)	-0.14 (-0.69)	-0.2 (-1.41)
EMU CPI	-0.06 (-0.53)	-0.26* (-1.66)	-0.41* (-1.93)	-0.11 (-0.67)	-0.16 (-0.9)	-0.14 (-1.06)
EMU GDP	-0.27 (-1.39)	-0.4 (-1.6)	-0.78** (-2.27)	-0.35 (-1.3)	-0.46 (-1.58)	-0.39* (-1.84)
FR Bus. Conf.	0.08 (0.66)	0.21 (1.38)	0.19 (0.9)	-0.09 (-0.56)	0.26 (1.49)	0.2 (1.54)
ECB Ref. Rate	0.27** (2.42)	0.53** (3.66)	0.32 (1.63)	0.09 (0.59)	-0.08 (-0.45)	0.35** (2.88)
China CPI YoY	0.07 (0.57)	0.03 (0.18)	-0.09 (-0.42)	-0.01 (-0.08)	0.09 (0.53)	0.06 (0.47)
China Ind. Prod.	0.06 (0.54)	-0.15 (-1.06)	-0.09 (-0.47)	-0.17 (-1.11)	-0.41** (-2.48)	-0.08 (-0.63)
China PMI	0.05 (0.31)	0.23 (1.18)	0.05 (0.2)	0.28 (1.34)	0.49** (2.19)	0.26 (1.6)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

TABLE 22.7 Estimation of the Impact of Economic News Over Selected Industrial Metals Depending on the NBER Phase of The Business Cycle

		Aluminium	Copper	Nickel	Zinc	Lead	GSCI Ind. Metals
REC	Intercept	-3.06** (-1.98)	-4.47** (-2.9)	0.45 (0.29)	-2.78* (-1.8)	5.06** (3.27)	-4.67** (-3.03)
	AR	0.01 (0.3)	0.03 (1.23)	0.02 (0.64)	0.01 (0.47)	0.03 (0.93)	0.02 (1.17)
	Non-Farm Payroll	0.9** (3.54)	0.78** (2.37)	0.47 (1.05)	0.76** (2.15)	0.22 (0.57)	0.76** (2.75)
	ISM	0.12 (0.54)	0.19 (0.69)	-0.48 (-1.24)	-0.37 (-1.22)	-0.02 (-0.05)	-0.02 (-0.09)
	Jobless Claims	-0.23 (-1.39)	-0.49** (-2.32)	-0.89** (-3.1)	-0.41* (-1.79)	-0.32 (-1.31)	-0.42** (-2.39)
	US CPI MoM	-0.21 (-1.16)	0.24 (1.04)	0.11 (0.35)	-0.12 (-0.47)	0.17 (0.61)	0.02 (0.08)
	US Retail Sales	0.22 (1.43)	-0.11 (-0.54)	-0.95** (-3.57)	0.04 (0.19)	-0.27 (-1.19)	-0.02 (-0.15)
	Fed Target Rate	0.05 (0.3)	0 (0.01)	0.31 (0.97)	0.06 (0.23)	0.37 (1.34)	0.1 (0.52)
	US GDP	0.63 (1.23)	1.46** (2.19)	0.36 (0.4)	1.38* (1.92)	1.46* (1.87)	1.1* (1.95)
	ZEW Eco. Sent.	-0.46* (-1.87)	-1.06** (-3.31)	-0.64 (-1.47)	0.51 (1.48)	-0.78** (-2.09)	-0.85** (-3.16)
	IFO Expectations	-0.62** (-2.7)	-1.12** (-3.78)	-0.92** (-2.28)	-0.39 (-1.22)	-0.47 (-1.35)	-0.82** (-3.28)
	EMU CPI	-0.14 (-0.59)	-0.29 (-0.93)	-0.63 (-1.51)	-0.24 (-0.73)	-0.57 (-1.59)	-0.14 (-0.54)
	EMU GDP	-0.72 (-1.12)	-0.7 (-0.85)	-1.95* (-1.72)	-2.39** (-2.67)	-3.28** (-3.38)	-1.05 (-1.5)
	FR Bus. Conf.	0.25 (1.27)	0.6** (2.36)	0.59* (1.68)	-0.19 (-0.7)	0.94** (3.17)	0.41* (1.89)
	ECB Ref. Rate	0.37** (2.59)	0.81** (4.42)	0.48* (1.92)	0.42** (2.11)	-0.17 (-0.79)	0.54** (3.53)
EXP	China CPI YoY	0.18 (0.47)	0.09 (0.19)	-0.48 (-0.7)	1.02* (1.88)	0.55 (0.94)	0.17 (0.4)
	China Ind. Prod.	0.28 (1.02)	-0.29 (-0.82)	0.25 (0.5)	-0.63 (-1.62)	-0.67 (-1.59)	-0.04 (-0.14)
	China PMI	0.16 (0.98)	0.41** (1.98)	0.24 (0.83)	0.39* (1.74)	0.61** (2.52)	0.29* (1.67)
	Non-Farm Payroll	0.11 (1.06)	0.04 (0.31)	0.1 (0.56)	0.1 (0.7)	0.06 (0.37)	0.09 (0.81)
	ISM	0.19* (1.81)	0.44** (3.19)	0.43** (2.28)	0.34** (2.26)	0.28* (1.77)	0.28** (2.38)
	Jobless Claims	-0.17 (-1.47)	-0.2 (-1.35)	-0.03 (-0.17)	-0.14 (-0.83)	-0.13 (-0.72)	-0.16 (-1.25)
	US CPI MoM	-0.07 (-0.62)	-0.05 (-0.33)	-0.46** (-2.28)	-0.1 (-0.64)	-0.03 (-0.17)	-0.1 (-0.8)
	US Retail Sales	0.1 (0.54)	0.05 (0.22)	0.03 (0.09)	-0.03 (-0.11)	-0.14 (-0.52)	0.03 (0.15)
	Fed Target Rate	0.06 (0.26)	-0.01 (-0.03)	-0.49 (-1.26)	0.05 (0.15)	0.07 (0.2)	0 (0.01)
	US GDP	0.05 (0.26)	0.22 (0.89)	-0.02 (-0.07)	0.37 (1.39)	-0.23 (-0.8)	0.12 (0.58)
	ZEW Eco. Sent.	-0.02 (-0.12)	-0.02 (-0.13)	-0.06 (-0.23)	0.06 (0.32)	-0.06 (-0.29)	0 (-0.02)
	IFO Expectations	0.03 (0.19)	0.14 (0.69)	0.04 (0.14)	0.2 (0.91)	-0.01 (-0.02)	0.09 (0.53)
	EMU CPI	-0.06 (-0.41)	-0.27 (-1.52)	-0.32 (-1.33)	-0.07 (-0.37)	-0.03 (-0.17)	-0.13 (-0.86)
	EMU GDP	-0.22 (-1.1)	-0.38 (-1.43)	-0.66* (-1.84)	-0.15 (-0.53)	-0.18 (-0.6)	-0.33 (-1.47)
	FR Bus. Conf.	0 (0.01)	0.03 (0.15)	-0.01 (-0.03)	-0.03 (-0.15)	-0.08 (-0.35)	0.11 (0.68)
	ECB Ref. Rate	0.1 (0.56)	0.07 (0.27)	0.12 (0.37)	-0.42* (-1.64)	0.05 (0.17)	0.03 (0.16)
	China CPI YoY	0.05 (0.39)	0.02 (0.11)	-0.02 (-0.11)	-0.11 (-0.66)	0.04 (0.23)	0.05 (0.35)
	China Ind. Prod.	0.02 (0.15)	-0.12 (-0.81)	-0.15 (-0.72)	-0.08 (-0.47)	-0.36** (-2)	-0.08 (-0.64)
	China PMI	-0.64 (-1.49)	-0.75 (-1.37)	-0.35 (-0.46)	-0.07 (-0.11)	0.02 (0.03)	0.35 (0.75)

Notes: REC stands for 'Recession' and EXP for 'Expansion' according to the NBER Business Cycle Dating Committee. T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

run, the aluminium price seems to increase when employment and production perspectives are improving, and to decrease when these conditions deteriorate. In [Table 22.7](#), we observe that recession periods lead to an heightened sensitivity of aluminium to the world business cycle: it decreases with declining payrolls, ECB rates but increases with German leadings. On the contrary, during expansion periods, the aluminium price seems to react weakly to economic news (only to the US ISM).

2. *Copper*: The case of copper is very similar to aluminium, with a large discrepancy in terms of market reaction to economic news over recessions and expansions. During the full sample, the price of copper increases with the ISM, the US unemployment and the ECB main refinancing rate, and decreases with the German leadings and the EMU CPI. During recession periods, the copper price drops with decreases in US leadings, US unemployment and US GDP, and with lowered ECB rate and Chinese PMI (surprisingly to the downside). The negative EMU GDP and German IFO lead to increases in the price of copper. This result may be explained by the fact that we focus on the NBER recession dates, whereas Europe has effectively entered the financial crisis during late 2008—i.e. almost one year after the USA. During expansion periods, the copper price is found to rise only with positive ISM surprises.
3. *Nickel*: During the full sample, the price of nickel displays a negative sensitivity to surprises in the US Retail Sales, Jobless Claims, the EMU CPI and the EMU GDP. This surprising pattern features clearly a counter-cyclical reaction to economic news. When considering recession periods, we uncover again a negative sensitivity of the nickel price to the German IFO, the US Retail Sales, Jobless Claims and the EMU GDP. This commodity records nonetheless a positive reaction to increases in the ECB main refinancing rate, and in the French business confidence survey. During expansion periods, the price of nickel reacts positively to unexpected increases in the US ISM, and negatively to unexpected increases in the US CPI and the EMU GDP. This finding underlines the potential complexity of the relationship between industrial metals and the economic newsflow.
4. *Zinc*: During the full sample, the price of zinc only reacts positively to surprises in the US GDP. During recession periods, this sensitivity to the US GDP appears again, even with an increased importance. Besides, the price of zinc displays strongly procyclical reactions to the ISM, the US unemployment, the US GDP, the ECB rate and the Chinese PMI. On the contrary, during expansion periods, the price of zinc exhibits two kinds of reactions: one positive to the US ISM, and one negative to the ECB main refinancing rate. Hence, zinc prices are more likely to rise when US surveys bring positive surprises, and when the ECB is in the process of lowering its main refinancing rate.
5. *Lead*: During the full sample, the lead price reacts positively to positive surprises in the Chinese PMI, and to negative surprises in the Chinese industrial production. When looking at the NBER Business Cycle phases, we find that this negative sensitivity to industrial production is only valid during expansion periods. This result may be related to the aggressive monetary policy conducted by the Public Bank of China over the post-2008 period. In addition, during expansionary phases, we find a positive sensitivity of the lead price to the US ISM. During recession periods, this commodity displays a positive sensitivity to the Chinese PMI—as in the full sample case—and to the US GDP, and a negative reaction to the EMU GDP.

To sum up, we find common patterns across industrial metals: most of them display a procyclical behaviour during recessions, a weak reaction to news during expansion periods and a significant sensitivity to Chinese news during recessions.

22.2.4.4 Energy

[Tables 22.8](#) and [22.9](#) present the estimation results of Equations (22.3) and (22.4) in the case of energy commodities.

TABLE 22.8 Estimation of the Impact of Economic News Over Selected Energy Markets

	WTI	Brent	Gasoil	Natural gas	Heating oil	GSCI energy
Intercept	-0.95 (-0.61)	-4.46** (-2.92)	0.25 (0.16)	-1.27 (-0.82)	-3.4** (-2.21)	-2.1 (-1.37)
AR	0.05 (1.39)	0.05 (1.62)	0.05 (1.5)	0.01 (0.2)	0.05 (1.43)	0.03 (1.15)
Non-Farm Payroll	0.15 (0.92)	0.09 (0.57)	0.19 (1.26)	0.31 (1.23)	0.16 (1.02)	0.12 (0.79)
ISM	0.2 (1.22)	0.04 (0.23)	0.09 (0.61)	0.29 (1.14)	0.13 (0.8)	0.09 (0.6)
Jobless Claims	-0.58** (-3.47)	-0.34** (-2.17)	0 (0.02)	0.16 (0.63)	-0.33** (-2.03)	-0.25* (-1.68)
US CPI MoM	-0.1 (-0.6)	-0.07 (-0.41)	-0.05 (-0.3)	0.02 (0.08)	0.01 (0.07)	-0.04 (-0.27)
US Retail Sales	-0.27 (-1.36)	-0.29 (-1.54)	-0.43** (-2.41)	-0.13 (-0.43)	-0.22 (-1.15)	-0.25 (-1.4)
Fed Target Rate	0 (0.01)	0.23 (1.02)	0.26 (1.22)	-0.21 (-0.57)	0.1 (0.42)	0.08 (0.36)
US GDP	0.21 (0.69)	0.1 (0.35)	0.04 (0.14)	0.24 (0.52)	0.19 (0.64)	0.22 (0.79)
ZEW Eco. Sent.	-0.19 (-0.93)	-0.13 (-0.65)	-0.08 (-0.41)	-0.39 (-1.23)	-0.27 (-1.31)	-0.21 (-1.15)
IFO Expectations	-0.06 (-0.24)	-0.07 (-0.3)	-0.25 (-1.23)	0.18 (0.51)	0.04 (0.18)	-0.03 (-0.15)
EMU CPI	-0.31 (-1.48)	-0.19 (-0.98)	-0.39** (-2.07)	0.25 (0.79)	-0.29 (-1.43)	-0.21 (-1.11)
EMU GDP	0.05 (0.14)	0.09 (0.27)	0.07 (0.23)	0.5 (0.96)	0.05 (0.15)	0.1 (0.32)
FR Bus. Conf.	0.16 (0.77)	0.22 (1.14)	0.15 (0.8)	0.44 (1.42)	0.21 (1.04)	0.19 (1.06)
ECB Ref. Rate	0.28 (1.43)	0.3* (1.64)	-0.06 (-0.33)	0.18 (0.61)	0.18 (0.96)	0.29* (1.66)
China CPI YoY	0.03 (0.13)	0.01 (0.07)	0.06 (0.36)	0.15 (0.48)	0.19 (1)	0.06 (0.31)
China Ind. Prod.	-0.09 (-0.49)	-0.1 (-0.53)	-0.1 (-0.6)	0 (-0.02)	-0.01 (-0.07)	-0.13 (-0.79)
China PMI	0.25 (0.97)	0.07 (0.27)	0.02 (0.07)	0.29 (0.74)	0.12 (0.46)	0.14 (0.6)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

The main findings may be summarized as follows:

1. *WTI and Brent*: crude oil prices show a rather weak long-term relationship to economic news. During the full sample, the only market mover on these markets is the US Jobless Claims, with a negative sign.* During recession periods, crude oil prices seem to be negatively related to US Jobless Claims, US Retail Sales, the German IFO, the German ZEW and the EMU CPI. In the meantime, these commodity prices also exhibit a positive reaction to increases in the ECB Refinancing Rate. This latter result illustrates that a rising ECB target rate is interpreted as a good news on crude oil markets, as it reveals that the world economy is roaring. During expansion periods, we cannot identify any significant market mover on these markets.
2. *Gasoil*: The full sample analysis reveals two market movers for the gasoil price: the US Retail Sales and the EMU CPI with negative signs. Hence, the gasoil market features an interesting counter-cyclical behaviour. During recession periods, we identify various market movers: the price of gasoil displays a negative reaction to the US Retail Sales, the ISM, German surveys, the EMU CPI and the Chinese industrial production. Besides, gasoil shows a positive reaction to the Chinese CPI and the US Non-Farm Payroll. During expansionary phases, we cannot relate the gasoil price to economic news at statistically significant levels. This latter result is conform to what has been found for crude oil prices.
3. *Natural Gas*: The full sample analysis reveals no reaction of the natural gas market to economic news. Expansion periods are characterized by one market mover: the US ISM with a procyclical pattern. During recession periods, the natural gas price shows a negative reaction to the German ZEW, and a positive reaction to the EMU CPI. Overall, natural gas seems to be the energy market with the weakest link to the economic newsflow, even when accounting for periods of recession.
4. *Heating Oil*: In line with the WTI and the Brent prices, the heating oil price reacts only to the ISM—positively—over the long run. During expansion periods, there is no reaction to economic news at statistically significant levels. During recession periods, the results are similar to crude oil prices. In addition, we uncover a positive sensitivity to the US Non-Farm Payroll.

* Note that the Brent price also displays a positive sensitivity to the ECB Refinancing Rate.

TABLE 22.9 Estimation of the Impact of Economic News Over Selected Energy Markets Depending on the NBER Phase of The Business Cycle

	WTI	Brent	Gasoil	Natural Gas	Heating Oil	GSCI Energy
Intercept	-0.65 (-0.42)	-4.05** (-2.64)	0.62 (0.4)	-1.25 (-0.81)	-3.3** (-2.13)	-1.78 (-1.16)
AR	0.05 (1.43)	0.05* (1.66)	0.05 (1.45)	0.02 (0.3)	0.05 (1.49)	0.04 (1.19)
REC Non-Farm Payroll	0.63 (1.44)	0.67 (1.59)	0.82** (2.07)	1.03 (1.52)	0.77* (1.78)	0.8** (2.03)
ISM	0.5 (1.3)	-0.02 (-0.04)	-0.59* (-1.72)	-0.65 (-1.12)	-0.36 (-0.98)	-0.22 (-0.65)
Jobless Claims	-1.31** (-4.63)	-0.71** (-2.62)	-0.05 (-0.21)	0.26 (0.61)	-0.61** (-2.23)	-0.58** (-2.3)
US CPI MoM	0.01 (0.04)	-0.09 (-0.31)	0.25 (0.89)	-0.26 (-0.53)	-0.04 (-0.15)	0 (0.02)
US Retail Sales	-0.48* (-1.82)	-0.51** (-2.01)	-0.59** (-2.49)	-0.11 (-0.26)	-0.38 (-1.49)	-0.41* (-1.75)
Fed Target Rate	-0.01 (-0.04)	0.26 (0.88)	0.23 (0.82)	-0.4 (-0.84)	0.04 (0.13)	-0.01 (-0.03)
US GDP	1.31 (1.46)	1.6* (1.87)	0.23 (0.28)	0.3 (0.22)	0.82 (0.93)	1.16 (1.45)
ZEW Eco. Sent.	-1.21** (-2.81)	-0.93** (-2.25)	-0.21 (-0.54)	-1.44** (-2.18)	-1.35** (-3.22)	-1.23** (-3.19)
IFO Expectations	-1.02** (-2.56)	-1.01** (-2.64)	-0.87** (-2.44)	-0.51 (-0.84)	-0.51 (-1.32)	-0.88** (-2.46)
EMU CPI	-0.73* (-1.77)	-0.94** (-2.37)	-0.7* (-1.88)	1.11* (1.75)	-0.94** (-2.34)	-0.73** (-1.96)
EMU GDP	0.56 (0.5)	0.16 (0.15)	-0.16 (-0.16)	2.44 (1.42)	0.12 (0.11)	0.71 (0.71)
FR Bus. Conf.	0.91** (2.66)	0.96** (2.93)	0.37 (1.22)	0.71 (1.36)	0.73** (2.18)	0.81** (2.64)
ECB Ref. Rate	0.55** (2.24)	0.54** (2.28)	-0.02 (-0.1)	0.25 (0.66)	0.29 (1.21)	0.48** (2.2)
China CPI YoY	-0.34 (-0.51)	0.26 (0.4)	1.33** (2.19)	0.12 (0.12)	0.6 (0.91)	-0.08 (-0.13)
China Ind. Prod.	-0.61 (-1.27)	-0.61 (-1.34)	-1.08** (-2.49)	-0.41 (-0.56)	-0.53 (-1.13)	-0.64 (-1.48)
China PMI	0.4 (1.42)	0.27 (1.02)	0.23 (0.93)	0.4 (0.93)	0.31 (1.15)	0.32 (1.28)
EXP Non-Farm Payroll	0.08 (0.46)	0 (-0.03)	0.07 (0.44)	0.17 (0.64)	0.06 (0.33)	0 (0.01)
ISM	0.16 (0.87)	0.07 (0.4)	0.27 (1.6)	0.52* (1.84)	0.27 (1.48)	0.18 (1.11)
Jobless Claims	-0.21 (-1.03)	-0.17 (-0.87)	0.02 (0.12)	0.08 (0.24)	-0.18 (-0.91)	-0.09 (-0.48)
US CPI MoM	-0.15 (-0.74)	-0.05 (-0.26)	-0.15 (-0.84)	0.15 (0.5)	0.05 (0.24)	-0.05 (-0.29)
US Retail Sales	0.02 (0.07)	0.01 (0.02)	-0.24 (-0.86)	-0.13 (-0.27)	0.01 (0.02)	-0.01 (-0.05)
Fed Target Rate	-0.02 (-0.06)	0.18 (0.48)	0.23 (0.65)	0.2 (0.34)	0.19 (0.5)	0.18 (0.51)
US GDP	0.08 (0.24)	-0.08 (-0.27)	0 (0.02)	0.21 (0.43)	0.12 (0.39)	0.1 (0.34)
ZEW Eco. Sent.	0.12 (0.5)	0.12 (0.54)	-0.02 (-0.09)	-0.09 (-0.24)	0.06 (0.28)	0.09 (0.44)
IFO Expectations	0.39 (1.38)	0.36 (1.36)	0.04 (0.17)	0.49 (1.15)	0.29 (1.05)	0.36 (1.43)
EMU CPI	-0.15 (-0.64)	0.04 (0.2)	-0.3 (-1.4)	-0.08 (-0.21)	-0.08 (-0.33)	-0.04 (-0.17)
EMU GDP	0 (-0.01)	0.08 (0.23)	0.09 (0.29)	0.31 (0.57)	0.04 (0.13)	0.04 (0.12)
FR Bus. Conf.	-0.21 (-0.86)	-0.13 (-0.56)	0.05 (0.24)	0.32 (0.83)	-0.06 (-0.23)	-0.1 (-0.47)
ECB Ref. Rate	-0.13 (-0.39)	-0.08 (-0.27)	-0.11 (-0.38)	0.04 (0.08)	0.01 (0.04)	-0.03 (-0.12)
China CPI YoY	0.06 (0.29)	-0.02 (-0.1)	-0.06 (-0.3)	0.12 (0.38)	0.14 (0.7)	0.06 (0.33)
China Ind. Prod.	-0.01 (-0.03)	0 (-0.01)	0.09 (0.48)	0.07 (0.22)	0.08 (0.4)	-0.05 (-0.26)
China PMI	0.41 (0.55)	-0.44 (-0.62)	-1.2* (-1.81)	-1.04 (-0.92)	-0.47 (-0.65)	-0.33 (-0.5)

Notes: REC stands for 'Recession' and EXP for 'Expansion' according to the NBER Business Cycle Dating Committee. T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

To sum up, energy markets seem to share a common pattern, i.e. a weak reaction to economic news during periods of economic expansion. This limited link during expansion periods might reflect the fact that energy commodities are influenced by local factors on the physical side of the market, as in the case of the natural gas.* During recession periods, the US unemployment seems to play a leading role in the evolution of these markets.

* We wish to thank a referee for highlighting this point.

22.2.4.5 Agricultural Commodities

In this section, we detail the results obtained with agricultural commodities in Tables 22.10–22.13.

The main empirical findings can be summarized as follows:

1. *Corn*: During the full sample, corn prices exhibit a negative relationship with the Fed Target Rate, the ZEW Economic Sentiment index and the Chinese PMI. According to this long-term

TABLE 22.10 Estimation of the Impact of Economic News Over Selected Agricultural Markets

	Corn	Wheat	Coffee	Sugar	Cocoa
Intercept	3.79** (2.48)	0.22 (0.14)	−0.91 (−0.59)	−0.13 (−0.08)	−0.6 (−0.39)
AR	0.03 (0.98)	0.02 (0.5)	0.01 (0.16)	0.01 (0.17)	0.01 (0.33)
Non-Farm Payroll	0.15 (1.12)	0.22 (1.53)	0 (0.02)	−0.04 (−0.24)	0.03 (0.24)
ISM	−0.13 (−0.96)	0.2 (1.35)	−0.1 (−0.56)	0.16 (1.01)	0.05 (0.37)
Jobless Claims	0.09 (0.71)	0.16 (1.11)	−0.15 (−0.84)	−0.1 (−0.62)	0.02 (0.12)
US CPI MoM	0.18 (1.34)	0.23 (1.56)	0.27 (1.47)	0.16 (0.96)	0.09 (0.63)
US Retail Sales	0.12 (0.73)	0.03 (0.19)	0.54** (2.51)	0.24 (1.23)	−0.04 (−0.25)
Fed Target Rate	−0.38* (−1.95)	−0.32 (−1.54)	−0.25 (−0.96)	−0.13 (−0.57)	0 (−0.01)
US GDP	0.19 (0.78)	0.18 (0.67)	0.3 (0.89)	0.13 (0.44)	−0.21 (−0.81)
ZEW Eco. Sent.	−0.41** (−2.46)	−0.29 (−1.62)	−0.44* (−1.94)	−0.32 (−1.57)	0.04 (0.23)
IFO Expectations	−0.14 (−0.75)	−0.04 (−0.18)	0.13 (0.51)	0.18 (0.81)	0.16 (0.81)
EMU CPI	−0.21 (−1.27)	−0.4** (−2.22)	−0.48** (−2.16)	−0.52** (−2.57)	−0.32* (−1.77)
EMU GDP	0.14 (0.53)	0.11 (0.38)	0.24 (0.65)	0.29 (0.88)	−0.17 (−0.59)
FR Bus. Conf.	0.11 (0.71)	0.31* (1.73)	0.22 (1.01)	0.32* (1.64)	0.31* (1.79)
ECB Ref. Rate	0.19 (1.19)	0.26 (1.54)	0.34 (1.6)	0.3 (1.59)	−0.15 (−0.92)
China CPI YoY	−0.22 (−1.35)	−0.24 (−1.36)	0.23 (1.09)	0.59** (3.03)	−0.28 (−1.6)
China Ind. Prod.	−0.05 (−0.35)	−0.04 (−0.25)	0.21 (1.02)	0.33* (1.8)	0.02 (0.11)
China PMI	−0.39* (−1.86)	−0.39* (−1.73)	0.26 (0.94)	0.41 (1.63)	−0.06 (−0.25)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

TABLE 22.11 Estimation of the Impact of Economic News Over Selected Agricultural Markets

	Cotton	Soybean	Rice	GSCI Agri.
Intercept	2.35 (1.53)	−0.87 (−0.57)	4.62** (3.01)	2.46 (1.61)
AR	0 (0.11)	0.02 (0.72)	0.02 (0.79)	−0.01 (−0.27)
Non-Farm Payroll	0.18 (1.28)	0.16 (1.39)	−0.01 (−0.05)	0.14 (1.56)
ISM	0.11 (0.77)	0.1 (0.86)	0.11 (0.88)	0.08 (0.84)
Jobless Claims	−0.12 (−0.89)	0.12 (1.07)	0 (−0.01)	0.05 (0.52)
US CPI MoM	0.04 (0.32)	0.22* (1.91)	0.19 (1.43)	0.18** (2)
US Retail Sales	0.22 (1.3)	0.01 (0.11)	0.08 (0.51)	0.09 (0.79)
Fed Target Rate	0.04 (0.19)	−0.12 (−0.72)	0.06 (0.29)	−0.23* (−1.75)
US GDP	0.18 (0.7)	0.11 (0.5)	0.24 (1.01)	0.15 (0.91)
ZEW Eco. Sent.	−0.21 (−1.22)	−0.32** (−2.18)	0.05 (0.33)	−0.31** (−2.76)
IFO Expectations	0 (0.01)	−0.09 (−0.54)	−0.1 (−0.59)	−0.06 (−0.5)
EMU CPI	−0.37** (−2.15)	−0.38** (−2.62)	−0.01 (−0.08)	−0.36** (−3.2)
EMU GDP	−0.41 (−1.46)	−0.05 (−0.2)	0.18 (0.68)	0.05 (0.26)
FR Bus. Conf.	0.33** (1.98)	0.13 (0.92)	0.25 (1.6)	0.19* (1.69)
ECB Ref. Rate	0.19 (1.16)	0.03 (0.19)	0.12 (0.8)	0.2* (1.88)
China CPI YoY	0.36** (2.18)	−0.12 (−0.87)	0.08 (0.5)	−0.08 (−0.74)
China Ind. Prod.	0.03 (0.19)	0.19 (1.46)	−0.05 (−0.32)	0.03 (0.24)
China PMI	0.04 (0.17)	−0.12 (−0.67)	−0.03 (−0.17)	−0.19 (−1.36)

Notes: T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

perspective, corn prices seem to be contra-cyclical—fearing for instance a tighter monetary policy from the Fed. These negative sensitivities are also valid during periods of recession. However, we also find a positive reaction to the US Non-Farm Payroll and the US CPI. To sum up, corn prices are found to increase when the Fed lowers its target rate, and when the US payroll and inflation figures are rising. During expansion periods, this commodity market does not seem to react statistically significantly to economic news.

2. *Wheat*: The price of wheat behaves similarly to corn, with two minor differences: (i) during recession periods, the wheat price does not react to the US payroll figures, and (ii) during expansion

TABLE 22.12 Estimation of the Impact of Economic News Over Selected Agricultural Markets Depending on the NBER Phase of the Business Cycle

		Corn	Wheat	Coffee	Sugar	Cocoa
REC	Intercept	4.01** (2.61)	0.62 (0.4)	-0.87 (-0.56)	-0.2 (-0.13)	-0.24 (-0.16)
	AR	0.03 (0.88)	0.01 (0.35)	0.01 (0.16)	0.01 (0.15)	0.01 (0.33)
	Non-Farm Payroll	0.66* (1.88)	0.37 (0.97)	0.31 (0.64)	0.72* (1.67)	0.7* (1.85)
	ISM	-0.03 (-0.09)	0.25 (0.74)	-0.32 (-0.79)	-0.16 (-0.43)	-0.08 (-0.24)
	Jobless Claims	-0.02 (-0.08)	0.28 (1.12)	0.03 (0.1)	0 (0.01)	-0.08 (-0.31)
	US CPI MoM	0.5** (1.97)	0.51* (1.86)	0.29 (0.85)	0.1 (0.32)	0.22 (0.8)
	US Retail Sales	0.13 (0.61)	0.14 (0.6)	0.88** (3.1)	0.25 (0.96)	0.08 (0.33)
	Fed Target Rate	-0.53** (-2.11)	-0.57** (-2.1)	-0.39 (-1.16)	-0.2 (-0.65)	0.22 (0.8)
	US GDP	0.79 (1.1)	0.67 (0.86)	0.42 (0.44)	0.02 (0.02)	0.1 (0.13)
	ZEW Eco. Sent.	-0.87** (-2.52)	-0.73* (-1.95)	-0.39 (-0.84)	-0.35 (-0.83)	-0.24 (-0.63)
	IFO Expectations	-1.1** (-3.44)	-1.19** (-3.42)	-0.14 (-0.32)	-0.36 (-0.92)	-0.16 (-0.47)
	EMU CPI	-0.53 (-1.59)	-0.9** (-2.5)	-0.95** (-2.12)	-0.59 (-1.46)	-0.82** (-2.29)
	EMU GDP	0.91 (1.02)	0.7 (0.72)	0.4 (0.33)	-0.05 (-0.05)	-1.39 (-1.44)
	FR Bus. Conf.	0.37 (1.33)	0.38 (1.28)	0.26 (0.7)	0.66** (1.98)	0.68** (2.31)
	ECB Ref. Rate	0.23 (1.17)	0.42** (1.96)	0.55** (2.05)	0.36 (1.51)	-0.34 (-1.62)
	China CPI YoY	-0.24 (-0.45)	-0.02 (-0.03)	0.59 (0.8)	-0.32 (-0.48)	-0.94 (-1.61)
	China Ind. Prod.	-0.62 (-1.61)	-0.67 (-1.59)	0.12 (0.23)	0.31 (0.66)	0.83** (1.98)
	China PMI	-0.45** (-2.01)	-0.5** (-2.04)	0.25 (0.82)	0.38 (1.39)	0 (0)
EXP	Non-Farm Payroll	0.07 (0.47)	0.19 (1.26)	-0.05 (-0.26)	-0.17 (-0.97)	-0.08 (-0.51)
	ISM	-0.14 (-0.92)	0.19 (1.18)	-0.04 (-0.22)	0.25 (1.41)	0.11 (0.66)
	Jobless Claims	0.16 (0.99)	0.11 (0.6)	-0.25 (-1.15)	-0.15 (-0.77)	0.09 (0.51)
	US CPI MoM	0.06 (0.36)	0.13 (0.74)	0.27 (1.23)	0.18 (0.91)	0.01 (0.04)
	US Retail Sales	0.1 (0.41)	-0.1 (-0.39)	0.08 (0.25)	0.24 (0.79)	-0.2 (-0.76)
	Fed Target Rate	-0.25 (-0.81)	-0.03 (-0.09)	-0.05 (-0.12)	0.02 (0.06)	-0.37 (-1.11)
	US GDP	0.1 (0.38)	0.09 (0.32)	0.28 (0.78)	0.15 (0.46)	-0.25 (-0.87)
	ZEW Eco. Sent.	-0.27 (-1.44)	-0.17 (-0.83)	-0.46* (-1.8)	-0.32 (-1.38)	0.13 (0.65)
	IFO Expectations	0.33 (1.46)	0.54** (2.19)	0.26 (0.84)	0.43 (1.57)	0.3 (1.25)
	EMU CPI	-0.08 (-0.43)	-0.21 (-1.01)	-0.32 (-1.24)	-0.49** (-2.09)	-0.14 (-0.7)
	EMU GDP	0.07 (0.24)	0.05 (0.17)	0.22 (0.58)	0.33 (0.95)	-0.05 (-0.16)
	FR Bus. Conf.	0.02 (0.12)	0.32 (1.45)	0.2 (0.74)	0.16 (0.66)	0.13 (0.63)
	ECB Ref. Rate	0.1 (0.41)	-0.04 (-0.15)	-0.06 (-0.16)	0.15 (0.47)	0.16 (0.57)
	China CPI YoY	-0.21 (-1.29)	-0.26 (-1.45)	0.19 (0.85)	0.68** (3.32)	-0.21 (-1.17)
	China Ind. Prod.	0.05 (0.31)	0.07 (0.42)	0.22 (1)	0.33 (1.63)	-0.14 (-0.77)
	China PMI	0.43 (0.73)	0.64 (1)	0.53 (0.66)	0.61 (0.84)	-0.07 (-0.11)

Notes: REC stands for 'Recession' and EXP for 'Expansion' according to the NBER Business Cycle Dating Committee. T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

TABLE 22.13 Estimation of the Impact of Economic News Over Selected Agricultural Markets Depending Over the NBER Cycles

		Cotton	Soybean	Rice	GSCI Agri.
REC	Intercept	2.34 (1.52)	-0.3 (-0.19)	4.65** (3.02)	2.99* (1.94)
	AR	0 (0.08)	0.01 (0.56)	0.02 (0.88)	-0.01 (-0.39)
	Non-Farm Payroll	0.45 (1.23)	0.31 (1.02)	-0.25 (-0.72)	0.45* (1.88)
	ISM	-0.17 (-0.54)	-0.14 (-0.53)	0.42 (1.41)	0 (0.01)
	Jobless Claims	-0.14 (-0.59)	0 (-0.01)	0.27 (1.21)	0.07 (0.43)
	US CPI MoM	-0.06 (-0.23)	0.27 (1.23)	0.26 (1.05)	0.37** (2.18)
	US Retail Sales	0.18 (0.82)	0.11 (0.59)	0.21 (1.01)	0.15 (1.06)
	Fed Target Rate	0.11 (0.44)	-0.27 (-1.23)	-0.08 (-0.32)	-0.4** (-2.36)
	US GDP	0.7 (0.95)	-0.04 (-0.07)	0.87 (1.25)	0.53 (1.08)
	ZEW Eco. Sent.	-0.48 (-1.34)	-0.76** (-2.52)	-0.52 (-1.56)	-0.7** (-2.96)
	IFO Expectations	-0.83** (-2.5)	-0.7** (-2.53)	-0.59* (-1.91)	-0.73** (-3.35)
	EMU CPI	-1.34** (-3.9)	-0.86** (-2.99)	0.19 (0.59)	-0.8** (-3.55)
	EMU GDP	-0.79 (-0.86)	-0.39 (-0.5)	0.24 (0.27)	0.24 (0.39)
	FR Bus. Conf.	0.76** (2.68)	0.38 (1.61)	0.75** (2.81)	0.4** (2.12)
	ECB Ref. Rate	0.43** (2.13)	0 (-0.01)	-0.03 (-0.18)	0.29** (2.14)
EXP	China CPI YoY	-0.03 (-0.06)	-0.36 (-0.77)	0.93* (1.78)	-0.16 (-0.43)
	China Ind. Prod.	0.38 (0.96)	-0.49 (-1.47)	-0.31 (-0.82)	-0.37 (-1.4)
	China PMI	0.07 (0.3)	0 (-0.02)	-0.18 (-0.83)	-0.22 (-1.46)
	Non-Farm Payroll	0.12 (0.84)	0.13 (1.06)	0.04 (0.27)	0.09 (0.91)
	ISM	0.19 (1.25)	0.17 (1.29)	0.03 (0.22)	0.11 (1.05)
	Jobless Claims	-0.11 (-0.68)	0.2 (1.39)	-0.12 (-0.78)	0.04 (0.4)
	US CPI MoM	0.08 (0.48)	0.21 (1.51)	0.18 (1.15)	0.11 (1.05)
	US Retail Sales	0.27 (1.08)	-0.1 (-0.48)	-0.08 (-0.35)	-0.01 (-0.04)
	Fed Target Rate	-0.04 (-0.13)	0.08 (0.31)	0.23 (0.78)	-0.06 (-0.28)
	US GDP	0.11 (0.42)	0.12 (0.54)	0.16 (0.63)	0.1 (0.56)
	ZEW Eco. Sent.	-0.13 (-0.65)	-0.19 (-1.17)	0.22 (1.21)	-0.21 (-1.6)
	IFO Expectations	0.39* (1.69)	0.2 (1.05)	0.11 (0.52)	0.25* (1.65)
	EMU CPI	-0.03 (-0.13)	-0.22 (-1.3)	-0.08 (-0.43)	-0.2 (-1.51)
	EMU GDP	-0.37 (-1.26)	-0.01 (-0.05)	0.17 (0.62)	0.03 (0.15)
	FR Bus. Conf.	0.13 (0.61)	0.02 (0.13)	0.02 (0.09)	0.1 (0.75)
	ECB Ref. Rate	-0.27 (-1.03)	0.08 (0.37)	0.3 (1.19)	0.03 (0.19)
	China CPI YoY	0.4** (2.31)	-0.11 (-0.77)	-0.02 (-0.1)	-0.07 (-0.65)
	China Ind. Prod.	-0.04 (-0.25)	0.31** (2.2)	0 (0.02)	0.09 (0.84)
	China PMI	0.56 (0.92)	-0.47 (-0.91)	0.48 (0.83)	0.36 (0.88)

Notes: REC stands for 'Recession' and EXP for 'Expansion' according to the NBER Business Cycle Dating Committee. T-stats are between brackets. Italic figures indicate statistically significant figures at a 10% risk level.

periods, this commodity price exhibits some sensitivity to the German IFO, i.e. to European economic conditions.

3. *Coffee*: The price of coffee is found to react positively to the US Retail Sales, and negatively to the German ZEW and the EMU CPI. By disentangling recession from expansion periods, we get the insight that the reaction to US Retail Sales and the EMU CPI occurs during recessions. On the contrary, the negative reaction to surprises in the German ZEW is a feature of expansion periods. Moreover, during recession periods, we uncover that the price of coffee exhibits a positive sensitivity to the ECB main refinancing rate.

4. *Sugar*: During the full sample, we find a positive reaction of the sugar price to the Chinese PMI and CPI, and to the French business confidence index. Conversely, our results highlight a negative reaction to the EMU CPI. During periods of recession, the sugar price shows a statistically significant and positive reaction to the US Non-Farm Payroll and to the French business confidence index. During expansion periods, this commodity price reacts negatively to positive surprises in the EMU CPI, and to negative surprises in Chinese CPI.
5. *Cocoa*: In the long term, the price of cocoa has a positive reaction to the French business confidence index, and a negative reaction to the EMU CPI. During recession periods, the cocoa price reacts positively to the US Non-Farm Payroll, to the French business confidence index and to the Chinese industrial production. However, we also record a negative reaction when it comes to positive surprises in the EMU CPI. During expansion periods, we cannot detect any statistically significant reaction of the cocoa price to economic news.
6. *Cotton*: The full sample analysis reveals that the cotton price reacts negatively to surprises in the EMU CPI, and positively to surprises in the French business confidence index and the Chinese CPI. During recession periods, the price of cotton shows a negative reaction to the German IFO and the EMU CPI. Besides, we uncover a positive reaction to the French business confidence index and to the ECB Refinancing Rate. During expansionary phases, the price of cotton exhibits a positive reaction to surprises in the IFO index and in the Chinese CPI index and the Chinese CPI. During recession periods, the price of cotton shows a negative reaction to the German IFO and the EMU CPI. Besides, we uncover a positive reaction to the French business confidence index and to the ECB Refinancing Rate. During expansionary phases, the price of cotton exhibits a positive reaction to surprises in the IFO index and in the Chinese CPI.
7. *Soybean*: During the full sample, we record a positive sensitivity of the soybean price to surprises in the US CPI. In addition, we detect a negative reaction to surprises in the German ZEW and the EMU CPI. During recession periods, the price of soybean exhibits a negative relationship with surprises in the German ZEW, the German IFO and the EMU CPI. During expansion periods, we find one market mover for this commodity price at statistically significant levels, i.e. the Chinese industrial production (with a positive sign).
8. *Rice*: During the full sample and expansionary phases, we are unable to detect any significant influence on the price of rice. During recession periods, we identify one negative relationship to the German IFO, and two positive reactions to the French business confidence index and the Chinese CPI.

Overall, agricultural products are characterized by a rather complicated pattern, as we find little similarities within that particular class of commodities. It seems that agricultural markets are more sensitive to local factors (e.g. geography and climate) than other commodities.* Finally, it is noteworthy to remark that we find almost no market mover for this type of commodity during expansionary phases (as in the case of energy markets).

During 1999–2011, can we identify an increasing sensitivity of commodity markets to the business cycle? As discussed in the next section, a positive reply would provide some empirical support to the view that commodities are increasingly correlated with financial markets.

22.2.5 Rolling Analysis

In what follows, we perform a rolling analysis, i.e. we run regressions similar to the previous section by using a rolling three-year sample. We are therefore able to compute the percentage of news having a statistical impact on commodity markets with a clearer view of how it evolves through time.

* We wish to thank a referee for stressing this fact.

In addition, for each of the estimates obtained, we compute the average absolute of β as detailed in Equation (22.3). This provides us with a sensitivity index of the reaction of commodities to economic news.

Formally, we compute the number of news whose β has a t -statistic greater than 1.64 (in absolute value). This corresponds to the quantile of a scaled Gaussian distribution at the 10% risk level. The sensitivity index for a given market—consistently with Equation (22.3)—is given by:

$$B = \frac{1}{I} \sum_{i=1}^I |\hat{\beta}_i|, \quad (22.5)$$

with B the sensitivity index, $\hat{\beta}_i$ the estimated sensitivities to the individual news and I the total number of news.

Figure 22.1 shows the evolution of the percentage of market movers for the GSCI sub-indices. Figure 22.2 contains the evolution of the sensitivity scores. Two main stylized facts can be listed:

1. Commodity markets exhibit a time-varying sensitivity to business cycle indicators. In Figures 22.1 and 22.2, the percentage of news to which commodity markets react and the cumulated sum of the absolute β vary strongly over time. For industrial metals and agricultural products, this time-varying pattern follows closely the phases of the US business cycle, according to the NBER. For energy markets, our results show a strong reaction to economic news during 2008–2009, especially what concerns the sensitivity score. These markets display a high sensitivity to Chinese news in 2008, and then to US news in 2009. The sensitivity to Chinese news is also observable in 2009 for industrial metals and agricultural products. Finally, precious metals are characterized by an increased sensitivity to news during 2004–2006. This period corresponds to rising inflation and tighter monetary policy. Under such circumstances, commodity markets typically display an increased sensitivity to economic signals. Precious metals also show an increased sensitivity to European news after the 2008–2009 crisis, i.e. their price tends to rise with the level of risk aversion.
2. The reaction of commodity markets to news depends on recession/expansion phases, but the picture looks globally more complex. In Figures 22.1 and 22.2, we notice that precious metals react to European news despite the end of the US economic crisis. The European sovereign crisis could




FIGURE 22.1 Percentage of news with a statistical impact over different sectors of the commodity market.




FIGURE 22.2 Rolling sensitivity index of news with a statistical impact over different sectors of the commodity market.

therefore have led to a rising level of risk aversion that triggered a surge in the price of precious metals (through a flight-to-quality event). A similar comment can be made during the 2004–2005 US slowdown in leading indicators: agricultural prices have shown an increased sensitivity to economic news, both in terms of number and average market beta.

Having detailed the reaction of each type of commodity to the economic newsflow, we investigate more closely in the next section how commodity prices vary along the business cycle in various geographical zones.

22.3 Economic Regimes and Commodity Markets as an Asset Class

This section is devoted to the understanding of the stylized behaviour of commodity prices depending on various phases of the business cycle. First, we present the structure of the Markov regime-switching model. Second, we investigate to which business cycle commodity markets seem to be mostly related. Third, we determine during which phase the commodity prices are more likely to rise or fall.

22.3.1 Measuring the Business Cycle

We recall here how to measure business cycles according to Hamilton (1989)'s methodology. Let X_t be a vector of variables representing the position of a given economy in the business cycle at time t . This economy is composed of N states, which are defined by a different expectation and covariance of X_t such that for the state $i \in [1 : N]$:

$$X_t \sim N(\mu_i, \Sigma_i), \quad (22.6)$$

with μ_i , the vector of expectations for the elements of X_t , and Σ_i , the covariance matrix. When X_t is conditionally Gaussian, its unconditional distribution is a mixture of Gaussian distributions, that can

encompass non-normality. The probability at each time t for the economy to be assigned to state i is time-varying. However, the transition matrix—i.e. the matrix containing the unconditional probabilities to move from state 1 to state 2—are fixed. Let us denote P this transition matrix. The parameters driving the model are the first two moments for each state, the probabilities involved in the transition matrix and the initial probabilities at time $t = 0$.

One of the main challenges for this model lies in the estimation of the parameters,* which are estimated by maximizing the log-likelihood function for a given time series. This log-likelihood function is given by:

$$L(\theta) = \sum_{t=1}^T \log f(X_t | \mathcal{F}_{t-1}, \theta), \quad (22.7)$$

with $f(\cdot)$ the conditional density of X_t , T the total number of observations and θ a vector containing the parameters to be estimated. $\mathcal{F}_{t-1}$ represents the information available at time t :

$$\mathcal{F}_{t-1} = \sigma\{X_{t-1}, X_{t-2}, \dots, X_1\}. \quad (22.8)$$

For a given θ , the conditional density at time t is:

$$f(X_t | \mathcal{F}_{t-1}, \theta) = \sum_{i=1}^N f(X_t | \mathcal{F}_{t-1}, \theta, S_t = i) P(S_t = i | \mathcal{F}_{t-1}), \quad (22.9)$$

with S_t the underlying state variable at time t . Intuitively, S_t is defined as the latent state of the economy, which is unknown to the econometrician but which can be inferred from the data by using a recursive filter in the spirit of the Kalman filter. Indeed, the central feature of the model is that the transition between regimes is governed by a Markov chain that assigns the probability of falling into the i^{th} regime, according to the unobserved state variable.

$P(S_t = i | \mathcal{F}_{t-1})$ is the probability to be assigned the state i at time t , conditionally upon the information available at time $t - 1$. The expression of $f(X_t | \mathcal{F}_{t-1}, \theta, S_t = i)$ is given by:

$$\begin{aligned} & f(X_t | \mathcal{F}_{t-1}, \theta, S_t = i) \\ &= \frac{1}{(2\pi)^{\frac{k}{2}} |\Sigma_i|} \exp \left\{ -\frac{1}{2} (X_t - \mu_i)^T \Sigma_i^{-1} (X_t - \mu_i) \right\}, \end{aligned} \quad (22.10)$$

with k the size of the vector X_t (i.e. the number of variables used to measure the phase of the business cycle). The parameters of the Markov regime-switching model are estimated by selecting θ such that:

$$\theta^* : \max_{\theta} L(\theta) \quad (22.11)$$

As mentioned above, the estimation process requires the use of a recursive filter—known as the Expectation Minimization (EM) algorithm—detailed in Hamilton (1989)'s original article. As is standard practice in the literature, we use the EM filter to obtain explicitly the parameters of the Markov regime-switching model for a given θ .

22.3.2 To Which Business Cycle are the Commodity Markets Related?

To our best knowledge, previous literature identifies the USA as the leading business cycle for commodity markets (Frankel and Hardouvelis 1985, Barnhart 1989). This view is justified by the fact that the US GDP is closely related to the worldwide business cycle, especially given the openness of this economy.

* See Hamilton (1996) for a detailed presentation.

Besides, most commodity prices are labelled in US dollar, and the economic and monetary regimes under which the USA evolve can arguably be of a primary importance for commodity investors.

This hypothesis needs to be verified empirically. That is why we propose to conduct a formal statistical analysis of the performance of commodity markets through seven different domestic regimes in the USA, Europe, Germany, Canada, Australia, Brazil and China. These business cycles have been selected on the ground of the potential interactions between these economies and commodity markets. For instance, Australia and Canada are well-known commodity producers, while China, Europe and the USA are mostly on the demand side of commodities. Brazil lies in between the two categories, as both a producer and a consumer. Taken together, these case studies will help us understanding whether there is any regional-specific behaviour at stake in the price development of commodity markets, given that our sample includes various regions covering North and South America, Europe, Asia and Oceania.

During 1993–2011, we have gathered for each of these countries seasonally adjusted measures of industrial production (IP) and CPI. We choose to relate the business cycle to these two series for several reasons. First, the industrial production is a well-known cyclical measure of the business cycle. Whenever production capacities are tense (idle), we can infer that the economy is in a boom (bust). Second, standard approaches exist to compute IP and CPI figures, which eases the comparability across countries. Third, IP and CPI are available with a monthly frequency, whereas the GDP is only accessible for quarterly data.

Hence, this higher number of observations will translate into more precise estimates. Taken together, those three reasons explain our choices when it comes to the data selection for the macroeconomic variables of interest.* This data-set is composed of monthly series.

We estimate the Markov regime-switching models as detailed in [Section 22.3.1](#). We set the number of regimes $N = 2$: Regime 1 is characterized as ‘expansion’ (i.e. production and prices are rising together), while regime 2 is referred to as ‘recession’ (i.e. prices record a slightly negative variation and production goes through a sharp drop). Estimation results are presented in Table 22.14 and in [Figure 22.3](#).

According to Table 22.14, during expansion periods, the US industrial production is expected to grow by 3.94%/month, while during recession periods it decreases by 7.27%/month. Similar comments arise for the different countries investigated here.† Another common feature of these different cycles consists

TABLE 22.14 Descriptive Statistics of the Domestic Cycles Estimated with the Markov Regime-Switching Model.

			USA	Europe	Germany	Australia	Canada	China	Brazil
Expansion	Ind. Prod.	μ	3.949	3.101	3.696	4.512	2.977	15.291	4.179
		σ	0.661	0.91	1.012	0.806	1.116	1.636	1.922
	Inflation	μ	2.624	2.127	1.785	3.756	2.242	10.506	5.881
		σ	0.278	0.297	0.343	0.476	0.41	1.819	0.747
	Stat.	Pr. of expansion	0.984	0.981	0.969	0.936	0.963	0.978	0.988
		Freq. (months)	192	163	148	134	196	138	171
		Freq. (%)	0.828	0.901	0.818	0.59	0.856	0.603	0.747
Recession	Ind. Prod.	μ	−7.265	−12.28	−4.522	0.153	−4.561	12.153	−0.495
		σ	1.806	3.675	4.864	0.856	1.895	1.364	4.418
	Inflation	μ	1.827	0.234	0.351	1.78	0.104	0.557	41.933
		σ	0.994	0.371	0.347	0.358	0.617	0.53	13.333
	Stat.	Pr. of recession	0.92	0.849	0.878	0.918	0.821	0.967	0.947
		Freq. (months)	40	18	33	93	33	91	58
		Freq. (%)	0.172	0.099	0.182	0.41	0.144	0.397	0.253

* Note that for Europe the sample starts in 1996.

† In Brazil, the recession was characterized by a surge in inflation during the so-called ‘Tequila crisis’ in the mid-1990s.




FIGURE 22.3 Business conditions during 1993–2011 estimated with the Markov regime-switching model.

in the duration of each regime: expansion periods typically last for a longer period of time than recessions. For instance, in Germany, the economy is expected to spend 18.2% of the sample in recession, with the remaining data points falling into expansionary phases.

Figure 22.3 shows the associated smoothed transition probabilities, along with the NBER business cycle reference dates (published by the NBER's Business Cycle Dating Committee*) in grey. Recessions start at the peak of a business cycle and end at the trough.

Estimates for the USA correspond most of the time to the NBER phases of the business cycle. In the rest of the world, the domestic cycles exhibit strong specificities. For instance, in 2008, the Eurozone has entered a recessionary regime later than the USA, and also exited later. In 2011, Germany has been impacted by a slowdown, without being followed by any other country.

Next, for each of these domestic cycles, we attempt to relate the performance of the S&P 500 (as a proxy of equity markets) compared to the GSCI and its sub-indices (Agriculture, Energy, Industrial Metals and Precious Metals). By doing so, we aim at evaluating how commodity markets behave along the business cycle, by explicitly taking into account country- and market-specific effects.

We are interested in the forecasting horizon that leads to the greatest discrimination between asset performances during expansionary and recessionary regimes. Thus, we need to find h such that excess returns—once adjusted from volatility—are as different as possible.

Let $SR_h^{(i)}(m)$ be the Sharpe ratio obtained during state m for the asset i when considering that commodity markets are ahead of the business cycle by h months. We select h such that:

$$\max_h \sum_{j=1}^N \sum_{k=1, k \neq j, k > j}^N (SR_h^{(i)}(m_j) - SR_h^{(i)}(m_k))^2, \quad (22.12)$$

* See more on the NBER Business Cycle Expansions and Contractions at <http://www.nber.org/cycles.html>

with m_j , the j^{th} , state of the economy, and N , the total number of states. With $N = 2$, we have:

$$\max_h (SR_h^{(i)}(m_1) - SR_h^{(i)}(m_2))^2. \quad (22.13)$$

This maximum spread between the Sharpe ratios can be used to measure the influence of a given phase of the business cycle on a given market. The results obtained are presented during the full sample (1993–2011) in Table 22.15, and during the subsample period of ‘commodity boom’ (2004–2011) in Table 22.16.

The following comments arise:

- During the full sample, the GSCI shows a strong connection to the business cycle located in the USA, Europe, Germany and Canada. The maximum score ranges from 0.23 in the USA to 0.42 in Europe. Therefore, we uncover that the US business cycle does not display the strongest

TABLE 22.15 Worldwide Business Cycles and Commodities Performances: 1993–2011

	GSCI Agriculture						
	AU	BR	CA	CH	EU	GE	US
Recession	0.28%	−0.18%	−1.53%	1.22%	−4.20%	−3.10%	−2.77%
Expansion	−0.17%	0.11%	0.27%	−0.76%	0.10%	0.40%	0.49%
Lead	1	5	7	1	8	7	1
Score max.	0.04	0.03	0.14	0.19	0.25	0.25	0.24
	GSCI Energy						
	AU	BR	CA	CH	EU	GE	US
Recession	1.15%	−1.85%	−3.56%	0.50%	−10.17%	−5.82%	−3.63%
Expansion	−0.19%	1.15%	1.08%	0.22%	1.72%	1.96%	1.08%
Lead	1	2	5	8	3	3	1
Score max.	0.07	0.16	0.2	0.03	0.51	0.37	0.22
	GSCI Industrial Metals						
	AU	BR	CA	CH	EU	GE	US
Recession	−0.22%	−0.94%	−2.68%	0.52%	−6.10%	−2.76%	−2.80%
Expansion	1.07%	1.10%	1.15%	0.59%	1.40%	1.42%	1.23%
Lead	1	3	5	1	6	5	6
Score max.	0.11	0.17	0.29	0	0.51	0.3	0.3
	GSCI Precious Metals						
	AU	BR	CA	CH	EU	GE	US
Recession	1.14%	0.98%	1.68%	1.04%	0.41%	1.74%	0.64%
Expansion	0.39%	0.64%	0.56%	0.46%	0.82%	0.63%	0.69%
Lead	8	1	1	6	5	2	8
Score max.	0.08	0.06	0.14	0.1	0.07	0.06	0.03
	GSCI						
	AU	BR	CA	CH	EU	GE	US
Recession	0.57%	−1.04%	−2.99%	0.70%	−5.89%	−3.02%	−2.16%
Expansion	0.55%	1.18%	1.19%	0.38%	1.45%	1.55%	1.10%
Lead	1	3	5	8	3	3	5
Score max.	0.01	0.18	0.27	0.06	0.42	0.31	0.23
	SP500						
	AU	BR	CA	CH	EU	GE	US
Recession	−0.14%	−0.39%	−1.48%	0.75%	−4.74%	−1.68%	−2.71%
Expansion	0.95%	0.79%	0.87%	0.34%	0.98%	0.87%	1.11%
Lead	1	6	3	5	5	5	6
Score max.	0.13	0.13	0.23	0.09	0.45	0.23	0.38

TABLE 22.16 Worldwide Business Cycles and Commodities Performances: 2004–2011

	GSCI Agriculture						
	AU	BR	CA	CH	EU	GE	US
Recession	–1.61%	–3.83%	–4.67%	–5.02%	–5.44%	–6.62%	–2.76%
Expansion	1.38%	0.90%	1.11%	0.80%	1.06%	1.05%	1.06%
Lead	8	8	4	1	3	3	2
Score max.	0.22	0.29	0.37	0.33	0.37	0.48	0.25
	GSCI Energy						
	AU	BR	CA	CH	EU	GE	US
Recession	–2.54%	–8.20%	–6.42%	–10.54%	–9.47%	–13.94%	–5.78%
Expansion	1.99%	1.42%	1.24%	1.34%	1.20%	1.66%	1.53%
Lead	8	8	3	1	1	3	2
Score max.	0.24	0.39	0.32	0.41	0.41	0.65	0.31
	GSCI Industrial Metals						
	AU	BR	CA	CH	EU	GE	US
Recession	–1.42%	–5.10%	–4.49%	–8.90%	–7.87%	–9.01%	–4.83%
Expansion	3.37%	2.61%	2.22%	2.44%	2.64%	2.43%	2.95%
Lead	8	8	8	1	3	4	5
Score max.	0.36	0.47	0.4	0.66	0.64	0.61	0.5
	GSCI Precious Metals						
	AU	BR	CA	CH	EU	GE	US
Recession	1.97%	2.06%	1.88%	2.91%	0.97%	1.67%	1.70%
Expansion	1.12%	1.33%	1.36%	1.22%	1.53%	1.42%	1.42%
Lead	1	1	1	8	8	1	7
Score max.	0.11	0	0.02	0.09	0.07	0.04	0.06
	GSCI						
	AU	BR	CA	CH	EU	GE	US
Recession	–1.43%	–5.47%	–4.58%	–8.92%	–5.87%	–10.38%	–3.40%
Expansion	2.88%	2.08%	2.04%	2.13%	1.98%	2.32%	2.11%
Lead	8	8	3	1	3	3	2
Score max.	0.33	0.42	0.37	0.53	0.4	0.67	0.34
	SP500						
	AU	BR	CA	CH	EU	GE	US
Recession	–0.85%	–2.10%	–2.78%	–5.96%	–4.95%	–5.06%	–3.51%
Expansion	1.28%	0.88%	0.99%	1.16%	1.25%	1.05%	1.44%
Lead	6	8	5	1	3	5	4
Score max.	0.28	0.27	0.38	0.66	0.54	0.51	0.51

relationship with commodity markets. This surprising finding is consistent across the GSCI sub-indices and during the sub-sample as well.

- The agricultural and energy sectors have the highest correlation with the business cycle measures. For instance, the maximum score is equal to 0.51 concerning the relationship between the GSCI energy-index and the European business cycle.
- During the full sample, industrial metals are not found to be the most cyclical sector. However, during 2004–2011, industrial metals are characterized by the strongest connection with the business cycle measures (with a maximum score equal to 0.66 in China). This result illustrates the influence of the sustained Chinese growth on the demand for raw materials in the industrial sector.

- In the long run, we find a limited evidence of any connection between commodity markets and emerging countries such as China and Brazil, or with Australia. The bulk of the long-term relationship between macroeconomic conditions and commodity markets seems to occur in Europe and North America. For agricultural products, however, notice that the maximum score (equal to 0.19) is reached in China.
- During 2004–2011, the influence of Brazil and China has been growing: for industrial metals, the maximum score is obtained with the Chinese business cycle and a one-month forward looking horizon. Similarly, the influence of the Brazilian business cycle is as strong as the USA's for the energy and industrial metals sectors (with a maximum score close to 0.4).
- Most of the commodities are found to be procyclical: during expansion periods, commodity prices record a positive momentum. On the contrary, recession periods are characterized by retreating prices. This behaviour is very similar to the S&P 500.
- Precious metals stand out as the only exception. First, they display on average the weakest relationship to any business cycle. Second, they record either negative or close-to-zero performances during periods of growth, and positive returns during periods of recession. This finding is consistent with the role of gold as a 'safe haven' where to store value during recessions (Baur and McDermott 2010, Beaudry *et al.* 2011), and as a hedge against inflation during expansion regimes.
- Finally, when the S&P 500 seems to be anticipating the business cycle by five to six months, commodities present a larger variety of forecasting horizons (ranging from 1 to 8). During the full sample, these horizons are closer to 3–5 when focusing only on the highest maximum score. Hence, it seems that commodity markets contain less anticipation from market agents than equity markets.

Having detailed the relationship between commodity prices and the business cycle in emerging and developed countries, we investigate the following question: during which phase of the business cycle are commodity prices more likely to rise or fall?

22.3.3 Commodity Performances Depending on the Nature of Each Economic Regime

In this section, we use a large data-set of US economic time series during 1984–2011. We select the following key economic indicators: the industrial production, the consumer price index, consumer good inventories, durable good inventories, the unemployment rate and the Fed Target Rate. Hence, we capture various characteristics of the US business cycle, going from inventory to inflation cycles. The data-set is composed of monthly series, and comes from Bloomberg.

Concerning the number of regimes necessary to model the US economy, previous literature focuses on three states (Sichel 1994, Boldin 1996, Clements and Krolzig 2003). As pointed by Hamilton (1996), the assumption that the process describing the data presents a given number N of regimes cannot be tested by using the usual likelihood ratio test, as specific regularity conditions are not fulfilled.*

To test explicitly that a Markov regime-switching model with two states is superior to another model with three states for instance, we use Vuong (1989)'s test based on the distributional goodness-of-fit that the model is able to provide. This approach has been applied recently to US equity and credit markets by Ielpo (2012). Formally, to compare a Markov regime-switching model with N_i states to another model with N_j states, we need to compute the following test statistics:

$$t_{N_i, N_j} = \frac{1}{T} \sum_{t=1}^T \left(\log f_{\hat{\theta}_{N_i}}(X_t | \mathcal{F}_{t-1}) - \log f_{\hat{\theta}_{N_j}}(X_t | \mathcal{F}_{t-1}) \right), \quad (22.14)$$

* Nonetheless, Hamilton (1996) presents a variety of tests to determine whether an additional state is required to model the dynamics of the data.

Under the null hypothesis that the Markov regime-switching model with N_i states provides an equivalent fit to the model with N_j states, this statistic is distributed as:

$$\frac{t_{N_i, N_j}}{\hat{\sigma}_T} \sqrt{T} \sim N(0, 1), \quad (22.15)$$

with σ_T the standard deviation associated to the statistic t_{N_i, N_j} . The alternative hypothesis is that the test statistic is different from 0, and that the forecasting ability of both models is not the same. The standard deviation $\hat{\sigma}_T$ is estimated by using the Newey–West long-run variance (HAC) estimator.

The best model is selected so that it is preferred to models with a lower number of states, but equivalent to models with a higher number of states. The resulting model specification should be parsimonious to avoid the problem of overfitting (Bradley and Jansen 2004), while being as consistent as possible with the joint distribution of returns.

When performing Vuong (1989)'s test, we find that $N = 5$ as shown in Table 22.17. We obtain two expansion regimes, two recession regimes and one 'stalling regime'. These regimes are displayed in Figure 22.4, with descriptive statistics given in Table 22.18. The following comments arise:

- Regimes 1 and 5 are *expansionary* regimes:
 - i. Regime 5 is characterized by a 3.6%/month increase in industrial production, rising inflation, receding unemployment, but still—on average—decreasing Fed's decision rates. Inventories are weakly evolving.
 - ii. Regime 1 is very similar to Regime 5. As we observe strongly building inventories, and an increasing Fed Target Rate, we label this regime as 'strong growth'.
- Regime 2 and 3 are *recessionary* regimes:
 - i. Regime 2 is characterized by a strong decline in industrial production, with unemployment rate rising rapidly and lowered Fed decision rates. By looking at Figure 22.4, we understand that this regime corresponds to the burst of an economic crisis.

TABLE 22.17 Vuong (1989)'s Test to Select the Appropriate Number of Regimes

	1 state	2 states	3 states	4 states	5 states	6 states
1 state		-7.4	-10.05	-11.4	-13.31	-12.7
2 states			-6.48	-10.62	-13.49	-13.3
3 states				-5.01	-10.22	-8.47
4 states					-5.16	-4.86
5 states						-1.72
6 states						




FIGURE 22.4 Business conditions during 1984–2011 estimated with the Markov regime-switching model.

TABLE 22.18 Descriptive Statistics of the Estimated Economic Regimes

		Strong exp.	Slowdown	Strong Crisis	Stalling	Medium Expans.
		Reg. 1	Reg. 2	Reg. 3	Reg. 4	Reg. 5
Indus. Prod.	μ	0.041	0	-0.076	0.021	0.036
	σ	0.009	0.017	0.024	0.011	0.007
CPI	μ	0.029	0.025	0.007	0.036	0.031
	σ	0.003	0.004	0.017	0.006	0.004
Good Inv.	μ	0.055	-0.065	-0.048	0.095	0.005
	σ	0.008	0.016	0.027	0.011	0.008
Dur. Good Inv.	μ	0.048	-0.079	-0.041	0.099	0.006
	σ	0.009	0.016	0.023	0.011	0.01
Unempl. rate	μ	-0.382	1.172	2.276	-0.307	-0.411
	σ	0.18	0.213	0.527	0.265	0.199
Fed TR	μ	0.971	-2.629	-1.379	-0.348	-0.399
	σ	0.407	0.565	0.513	1.042	0.358
Freq. (months)		92	40	27	67	102
Freq. (%)		0.28	0.122	0.082	0.204	0.311

TABLE 22.19 Estimated Transition Matrix Between Regimes

	Reg. 1 (%)	Reg. 2 (%)	Reg. 3 (%)	Reg. 4 (%)	Reg. 5 (%)
Reg. 1	89	1	0	6	4
Reg. 2	7	93	0	0	0
Reg. 3	0	7	82	0	11
Reg. 4	2	0	3	91	3
Reg. 5	6	0	3	0	91

- ii. Regime 3 is a typical slowdown period: industrial production decreases—but not as much as for Regime 2—and the Fed decreases even more aggressively its decision rate, so as to limit as much as possible the building of unemployment. When examining the transition matrix presented in Table 22.19, the persistence of this regime is weaker than the persistence of Regime 2.
- Regime 4 is a ‘stalling’ regime, i.e. between expansion and recession. Industrial production grows by 2%/month, whereas good and durable good inventories are building up sharply, underlining the weak underlying consumers’ demand.

Table 22.20 presents the optimal forecasting horizon that should be retained to link the five regimes to the performances of commodities. We use the same metric as presented in Equation (22.13) to which we refer as the ‘maximum score’. Based on these results, we can describe the stylized performances of commodity sectors, depending on the phase of the business cycle.

Table 22.21 presents the performance by regime for each of the assets considered. The main findings can be summarized as follows:

- When entering the crisis, all commodities deliver strongly negative returns, except precious metals (with an increase by 1.29%/month on average). Hence, we verify the role of ‘safe haven’ for precious metals when entering a recession. Other commodity markets logically underperform compared to the S&P 500. For instance, the energy sector records a drop by -4.14%/month vs. -2.43%/month for equities.

TABLE 22.20 Optimal Forecasting Horizon for the Various Asset Classes

	GSCI Agri.	GSCI Energy	GSCI Ind. Metals	GSCI Prec. Metals	GSCI Total	SP500	10Y US	US Dollar
Leading +0	0.16	0.25	0.2	0.07	0.24	0.12	0.09	0.14
Leading +1	0.14	0.27	0.2	0.11	0.29	0.21	0.05	0.08
Leading +2	0.26	0.27	0.24	0.11	0.3	0.24	0.09	0.06
Leading +3	0.25	0.26	0.22	0.13	0.29	0.24	0.08	0.1
Leading +4	0.31	0.3	0.32	0.14	0.32	0.25	0.08	0.06
Leading +5	0.22	0.26	0.35	0.18	0.3	0.29	0.16	0.07
Leading +6	0.19	0.21	0.31	0.11	0.26	0.31	0.17	0.1
Leading +7	0.17	0.18	0.29	0.21	0.26	0.36	0.22	0.1
Leading +8	0.22	0.19	0.36	0.15	0.28	0.41	0.3	0.13
Leading +9	0.19	0.13	0.37	0.14	0.24	0.4	0.24	0.11
Leading +10	0.11	0.05	0.34	0.08	0.18	0.39	0.18	0.15
Leading +11	0.17	0.12	0.34	0.18	0.22	0.36	0.33	0.18
Leading +12	0.27	0.16	0.27	0.2	0.27	0.37	0.31	0.27

TABLE 22.21 Performance of Various Indices Depending on Each of the Economic Regimes

	Strong exp.	Slowdown	Strong Crisis	Stalling	Medium Expans.
	State 1 (%)	State 2 (%)	State 3 (%)	State 4 (%)	State 5 (%)
GSCI Agri.	0.47	-0.72	-2.87	1.27	0.27
GSCI Energy	2.63	0.32	-4.14	1.60	-0.18
GSCI Ind. Metals	2.11	0.21	-2.12	3.44	-0.45
GSCI Prec. Metals	0.55	-0.51	1.29	1.25	0.05
GSCI Total	1.76	-0.34	-2.63	1.36	-0.37
SP500	1.37	-0.80	-2.43	0.84	1.20
10Y US	1.36	-0.91	-2.58	-0.03	-1.66
US Dollar	-0.31	0.11	-0.29	0.24	-0.30

- During the prolonged slowdown period, central banks are stimulating the economy. As a consequence, we find that energy and industrial metals commodities are improving on average by, respectively, 0.32%/month and 0.21%/month. Other cyclical commodities exhibit decreasing prices.
- During the medium expansion regime, agricultural prices are rising by 0.27%/month. On the contrary, energy and industrial metals prices are decreasing by, respectively, -0.18%/month and -0.45%/month. More importantly, the S&P 500 outperforms all commodities. Investors would therefore not increase their exposure to commodity markets during periods of medium growth.
- During strong expansions, energy and industrial metals prices show positive returns, and even outperform equity returns. This strong expansionary phase is typically a period of 'commodity boom' for investors.
- During the stalling regime, all commodity prices are rising and outperform the returns on the S&P 500. With roaring inflation, even the precious metals deliver positive returns, consistent with the idea that these metals are also used as a hedge against inflation. Industrial metals are found to outperform all commodity indices.

By decomposing the business cycle into five regimes, we have successfully shown in this section that (long) investors should invest into commodities especially during the stalling regime, i.e. at the end of

the economic cycle (before the recession). According to our estimates, the stalling regime corresponds to 20% of the data sample.

22.4 Conclusion

The contributions of the paper are twofold: (i) evaluating in a regression framework with EGARCH effects the influence of economic news on each type of commodity (agricultural products, energy markets, industrial and precious metals), and (ii) studying the evolution of commodity prices along the business cycle in different geographic regions based on the class of Markov regime-switching models.

From an investment point of view, the main results may be summarized as follows. First, investors should not add commodities to their portfolio on the ground of a low correlation of the commodity with standard financial assets (such as bonds and equities). Indeed, we uncover a strong correlation of commodity markets with risky assets during economic downturns. Second, the economic influences on commodity prices are rather complex: the patterns detected for each type of commodity sub-index vary greatly between the geographical zones in Europe, the US, China, etc. Third, there is a cyclical rotation amongst commodity markets: during periods of strong growth, investors should overweigh industrial metals and energy vs. agricultural products and industrial metals. Finally, 'stalling' economic regimes correspond to periods during which commodities outperform the S&P 500.

References

- Alizadeh, A.H., Nomikos, N.K. and Poulialis, P.K., A Markov regime switching approach for hedging energy commodities. *J. Bank. Financ.*, 2008, **32**(9), 1970–1983.
- Andersen, L., Markov models for commodity futures: Theory and practice. *Quant. Finance*, 2010, **10**(8), 831–854.
- Barnhart, S.W., The effects of macroeconomic announcements on commodity prices. *Am. J. Agr. Econ.*, 1989, **17**(2), 389–403.
- Baur, D.G. and McDermott, T.K., Is gold a safe haven? International evidence. *J. Bank. Financ.*, 2010, **34**, 1886–1898.
- Beaudry, P., Collard, F. and Portier, F., Gold rush fever in business cycles. *J. Monetary Econ.*, 2011, **58**, 84–97.
- Boldin, M.D., A check on the robustness of Hamilton's Markov Switching model approach to the economic analysis of the business cycle. *Stud. Nonlinear Dyn. E.*, 1996, **1**, 35–46.
- Bradley, M.D. and Jansen, D.W., Forecasting with a nonlinear dynamic model of stock returns and industrial production. *Int. J. Forecasting*, 2004, **20**(2), 321–342.
- Caballero, R.J., Farhi, E. and Gourinchas, P.O., Financial crash, commodity prices, and global imbalances. *Brookings Pap. Eco. Act.*, 2008, Fall **2008**, 1–54.
- Chan, K.F., Treepongkaruna, S., Brooks, R. and Gray, S., Asset market linkages: Evidence from financial, commodity and real estate assets. *J. Bank. Financ.*, 2011, **35**(6), 1415–1426.
- Chevallier, J., Global imbalances, cross-market linkages, and the financial crisis: A multivariate Markov-switching analysis. *Econ. Modell.*, 2012, **29**(3), 943–973.
- Chng, M.T., Economic linkages across commodity futures: Hedging and trading implications. *J. Bank. Financ.*, 2009, **33**(5), 958–970.
- Clements, M.P. and Krolzig, H.M., Business cycle asymmetries: Characterization and testing based on Markov-switching autoregressions. *J. Bus. Econ. Stat.*, 2003, **21**(1), 196–211.
- Daskalaki, C. and Skiadopoulos, G., Should investors include commodities in their portfolios after all? New evidence. *J. Bank. Financ.*, 2011, **35**(10), 2606–2626.
- Dionne, G., Gauthier, G., Hammami, K., Maurice, M. and Simonato, J.G., A reduced form model of default spreads with Markov switching macroeconomic factors. *J. Bank. Financ.*, 2011, **35**, 1984–2000.

- Elder, J., Miao, H. and Ramchander, S., Impact of macroeconomic news on metal futures. *J. Bank. Financ.*, 2012, **36**(1), 51–65.
- Frankel, J.A. and Hardouvelis, G.A., Commodity prices, money surprises and fed credibility. *J. Money Credit Bank.*, 1985, **17**, 425–438.
- Geman, H. and Ohana, S., Time-consistency in managing a commodity portfolio: A dynamic risk measure approach. *J. Bank. Financ.*, 2008, **32**(10), 1991–2005.
- Hamilton, J., A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 1989, **57**(2), 357–384.
- Hamilton, J., *Time Series Analysis*, 2nd ed., 1996 (Princeton University Press: Princeton, NJ).
- Ielpo, F., Equity, credit and the business cycle. *Appl. Financ. Econ.*, 2012, **22**(12), 939–954.
- Miffre, J. and Rallis, G., Momentum strategies in commodity futures markets. *J. Bank. Financ.*, 2007, **31**(6), 1863–1886.
- Miltersen, K.L., Commodity price modelling that matches current observables: A new approach. *Quant. Finance*, 2003, **3**(1), 51–58.
- Sichel, D., Inventories and the three phases of the business cycle. *J. Bus. Eco. Stat.*, 1994, **12**, 269–277.
- Tang, K., Time-varying long-run mean of commodity prices and the modeling of futures term structures. *Quant. Finance*, 2012, **12**, 781–790.
- Tang, K. and Xiong, W., Index investing and the financialization of commodities. NBER Working Paper #16385, Cambridge, MA, 2010.
- Vuong, Q.H., Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 1989, **57**(2), 307–333.

Appendix: Database of News

The Appendix provides background information on the Bloomberg data-set of news used in the paper.

- *Non-Farm Payroll*: number of jobs added or lost in the US economy over the last month, not including jobs relating to the farming industry. It includes goods-producing, construction and manufacturing companies.
- *ISM*: survey from the Institute for Supply Management based on comments from purchasing managers in the manufacturing sector. It provides the earliest clues of how the US economy has fared during the previous four weeks.
- *Jobless Claims*: the weekly number of new jobless applications in the US.
- *US CPI MoM*: Consumer Price Index Month on Month.
- *US Retail Sales*: Retail and Food Services Sales in million US Dollars.
- *Fed Target Rate*: interest rate at which depository institutions actively trade balances held at the Federal Reserve.
- *US GDP*: US Gross Domestic Product.
- *ZEW Eco. Sent.*: survey from the Mannheim Centre for European Economic Research (ZEW). Expectations from 350 financial experts for the Euro-zone, Japan, Great Britain and the USA.
- *IFO Expectations*: survey from the Munich Society for the Promotion of Economic Research (IFO). 7000 monthly survey responses from firms in manufacturing, construction, wholesaling and retailing. Firms are asked to give their assessments of the current business situation and their expectations for the next six months.
- *EMU CPI*: Euro Area Consumer Price Index from Eurostat.
- *EMU GDP*: Euro Area Gross Domestic Product from Eurostat.
- *FR Bus. Conf.*: France Business Confidence index from INSEE.

- *ECB Ref. Rate*: minimum refinancing rate from the European Central Bank.
- *China CPI YoY*: China Consumer Price Index Year on Year
- *China Ind. Prod.*: China Industrial Production index.
- *China PMI*: China Purchasing Managers' Index.

Overall, the news cover different aspects of the business cycle:

- **real activity** with macroeconomic data (such as industrial production) or surveys,
- **inflation** dynamics with the CPI,
- **monetary** variables with the Fed Target Rate and the ECB's minimum refinancing rate.

The Dynamics of Commodity Prices

23.1	Introduction	501
23.2	Models and Estimation	503
	Spot Price Models • Estimation Approach	
23.3	Data and Empirical Results.....	506
	Data • Univariate Analysis • Cross-Commodity Market Analysis • Commodity and Equity Markets	
23.4	Economic Implications	515
	Options Valuation • Options Hedging	
23.5	Conclusions.....	519
	Acknowledgements	520
	References.....	520
	Appendix A: MCMC Estimation Details	522

Chris Brooks

Marcel Prokopczuk

In this paper we study the stochastic behavior of the prices and volatilities of a sample of six of the most important commodity markets and we compare these properties with those of the equity market. We observe a substantial degree of heterogeneity in the behavior of the series. Our findings show that it is inappropriate to treat different kinds of commodities as a single asset class as is frequently the case in the academic literature and in the industry. We demonstrate that commodities can be a useful diversifier of equity volatility as well as equity returns. Options pricing and hedging applications exemplify the economic impacts of the differences across commodities and between model specifications.

Keywords: Commodity prices; Stochastic volatility; Jumps; Markov chain Monte Carlo

JEL Classification: G1, G10, C3, C32

23.1 Introduction

Following the seminal work of Samuelson (1965), it is now widely accepted that commodity prices fluctuate randomly. Understanding the nature of this stochastic behavior is of crucial importance for decision makers engaging in commodity markets. Yet traditionally, with the probable exception of gold, commodities have not been in the focus of investors. However, interest has grown enormously over the last decade for a variety of reasons. First, following the relatively poor performances of stocks and Treasuries, investors have sought previously unexplored asset classes as potential new sources of returns. Second, the low correlation of commodity returns with equities and their ability to provide a hedge against inflation make them useful additions to portfolios. The liberalization of numerous markets has also increased corporates' requirements for hedging. This increased investment and hedging interest has led to a fast growth of the commodity derivatives market.

The aim of this paper is to study the stochastic behavior of commodity prices, both from an individual perspective but also concerning the cross-market linkages between commodities, and between commodities and the equity market. As such, this is the first piece of research to comprehensively apply a range of stochastic volatility models to commodities from several market segments. A good understanding of commodity prices' behavior and their interdependences as well as their relation to the equity market is important for investors, producers, consumers, and also policymakers.

We study the following issues for six major commodity markets. First, we investigate whether volatility does indeed behave stochastically and we estimate several models for the returns and volatility processes. Second, we examine the volatility of volatility, its persistence, and whether changes in prices and volatility are correlated. We additionally allow for jumps in both prices and volatility. We also investigate the linkages between the commodity markets considered. We determine whether volatilities are correlated across markets and whether the prices or the volatilities of different commodities jump at the same time. Moreover, we analyse the same questions for the linkages between each commodity and the equity market.

Our main findings are as follows. First, we find that, within the stochastic volatility framework, the models that allow for jumps provide a considerably better fit to the data than those that do not, although there is little to choose between the models allowing for jumps in returns only and those allowing for jumps in both returns and volatility. Second, we observe alternate signs in the relationships between returns and volatilities for different commodities—negative for crude oil and equities, close to zero for gasoline and wheat, and positive for gold, silver, and soybeans. We attribute these differences to variations in the relative balances of speculators and hedgers across the markets. We also find evidence of considerable differences in both the intensity and frequency of jumps, although all commodities are found to exhibit more frequent jumps than the S&P 500. We conclude that commodities have very different stochastic properties, and therefore that it is suboptimal to consider them as a single, unified asset class. To analyse the economic implications of the differences across commodities and between model specifications, as exemplars we employ applications focused on options pricing and hedging.

The stochastic behavior of prices, equity markets, and especially equity index markets, has received a great deal of attention. The non-normality of equity returns has been documented extensively. Motivated by the poor performance of the Black–Scholes–Merton options pricing formula that produced the well-known smile phenomenon, researchers have extended the simple Brownian motion in various directions. Merton (1976) was probably the first to suggest adding a discontinuous jump component to the continuous Brownian price process, and Ball and Torous (1985) empirically confirmed the merits of this approach. Scott (1987) and Heston (1993) suggested modeling volatility stochastically, and the latter study is able to derive a semi-closed-form solution for European option prices in the environment where volatility is stochastic. The two ideas were combined by Bates (1996) and Bakshi *et al.* (1997), who proposed employing stochastic volatility and price jumps to improve the description of the asset's price process with the aim of enhancing the accuracy of options pricing. Finally, Duffie *et al.* (2000) developed a general affine framework for asset prices, and suggested a model with stochastic volatility, price jumps, and, additionally, jumps in volatility. Andersen *et al.* (2002), Chernov *et al.* (2003), Eraker *et al.* (2003), and Eraker (2004) studied various versions of stochastic volatility models, with the latter two concluding that jumps in prices and volatility improve the description of S&P 500 price dynamics as well as the pricing of options written on the index. Asgharian and Bengtsson (2006) used this class of models to study the cross-market dependencies of various equity index markets.

Compared with this array of research on equity markets, state-of-the-art empirical studies on commodity markets are sparse. Of the few examples there are of such studies, Brennan and Schwartz (1985), Gibson and Schwartz (1990), Schwartz (1977), and Schwartz and Smith (2000) study the ability of continuous-time Gaussian factor models with constant volatility to describe the stochastic behavior

of the futures curve, mostly considering the crude oil market. Also studying the crude oil market only, Larsson and Nossman (2011) examine the empirical performance of several stochastic volatility models with jumps. Sorensen (2002) and Manoliu and Tompaidis (2002) apply the model of Schwartz and Smith (2000) to the agricultural and natural gas futures markets, respectively. Casassus and Collin-Dufresne (2005) mainly study the nature of risk premia in four different commodity markets; they also allow for the possibility of discontinuous price jumps, which Aiube *et al.* (2008) conclude are important in the crude oil market. A three-factor model incorporating prices, interest rates and the convenience yield is constructed by Liu and Tang (2011) to capture the stochastic relationship between the convenience yield level and its volatility for industrial commodities. Tang (2012) develops another three-factor model with a stochastic long-run mean. The possibility of stochastic volatility in commodities is considered by Geman and Nguyen (2005) for the soybean market, and by Trolle and Schwartz (2009) for crude oil.*

Papers considering multiple commodity markets and their dependencies rely mainly on simple return analyses. For example, Erb and Harvey (2006) and Gorton and Rouwenhorst (2006) study the benefits of investing in commodity markets; Kat and Oomen (2007a,b) conduct statistical analyses of commodity returns. In a paper related to ours, Du *et al.* (2011) estimate a bivariate stochastic volatility model for crude oil and two agricultural markets. However, the focus of their analysis is on pure volatility spillovers and their model does not allow for jumps in prices nor in volatility.

The remainder of the chapter is structured as follows. Section 23.2 describes the models estimated and the procedures employed to obtain the parameters, while Section 23.3 presents and discusses the data and the empirical results. Section 23.4 analyses the economic implications of employing different models in terms of options valuation and hedging errors. Finally, Section 23.5 concludes, while further details of the Markov chain Monte Carlo procedure are presented in an appendix.

23.2 Models and Estimation

23.2.1 Spot price models

The first model specification employed includes stochastic volatility only, and is therefore denoted SV. The log spot price $Y_t = \log S_t$ is assumed to follow the dynamics

$$dY_t = \mu(\tau)dt + \sqrt{V_t}dW_t^Y, \quad (23.1)$$

where W_t^Y is a standard Brownian motion.† To capture potential seasonal effects in the price dynamics, we assume a trigonometric function for the drift component in (23.1)

$$\mu(\tau) = \bar{\mu} + \eta \sin(2\pi(\tau + \zeta)), \quad (23.2)$$

where $\tau \in [0, 1]$ denotes the time fraction of the year elapsed. The average drift rate is denoted by $\bar{\mu}$. The parameter $\eta > 0$ controls the amplitude of the function and therefore captures the strength of the

* There is an enormous number of papers that model time-varying volatilities and correlations within a discrete-time GARCH-type framework. In the commodities area, many of these are on the oil market and focus on the objective of determining effective hedge ratios. A full survey of such work is beyond the scope of this paper, but relevant studies include Serletis (1994), Ng and Pirrong (1996), Haigh and Holt (2002), Pindyck (2004), Sadorsky (2006), Alizadeh *et al.* (2008), and Wang *et al.* (2008).

† Alternatively, one could specify the spot price dynamics as a mean-reverting process. We have done this, however the empirical results were inferior to the non-stationary Brownian motion.

seasonal effect, whereas ζ governs the periodicity of the process' drift capturing the form of the seasonality.* For markets that do not show any seasonal behavior we set $\eta = 0$, yielding a constant drift. For the volatility $\sqrt{V_t}$, we consider the square-root process

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^V, \quad (23.3)$$

where W_t^V is a second Brownian motion with $dW_t^Y dW_t^V = \rho dt$. The variance process is mean-reverting towards the long-run mean θ with speed κ . The parameter σ captures the volatility of volatility (vol-of-vol) and the kurtosis of returns increases for higher values of σ . The correlation between returns and volatility captures the skewness of the returns distribution. Positive values imply a skew to the right, and negative values a skew to the left.

This specification guarantees the positiveness of price and volatility at all times. Except for the seasonal adjustment, it is identical to the model proposed by Heston (1993).

Empirical evidence suggests that a pure continuous specification of the spot price cannot capture all salient features observed in financial data, and in particular the possibility of rapid price movements, i.e. jumps. To allow for jumps in the price dynamics, we follow Bates (1996) and add a Poisson process N_t^Y to specification (23.1), yielding the SVJ model

$$dY_t = \mu(\tau)dt + \sqrt{V_t}dW_t^Y + Z_t dN_t^Y. \quad (23.4)$$

The intensity λ_Y of N_t^Y is assumed to be constant and the jump sizes are assumed to be generated by a normal distribution, i.e. $Z_t \sim N(\mu_Y, \sigma_Y^2)$. Assuming that jumps occur infrequently but are relatively large (which is the natural perception of jumps as opposed to the diffusion components in prices), the jump component will mainly affect the tails of the return distribution. The specifications for the variance process and the seasonality adjustment remain unchanged.

Bakshi *et al.* (1997) and Bates (2000) find that although the inclusion of jumps in returns helps to describe the behavior of equity prices and the pricing of options, the model is still severely misspecified. Jumps in returns are transient, however, and hence a more persistent component is needed. Therefore, Duffie *et al.* (2000) introduce a model specification allowing for jumps in both returns and volatility. The return process remains identical to (23.4), but the variance process is amended by the jump process N_t^V , i.e.

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^V + C_t dN_t^V. \quad (23.5)$$

We assume that prices and volatility jump simultaneously, i.e. $N_t^V = N_t^Y$, which is commonly denoted as a SVCJ model. This assumption is motivated by the idea that periods of stress when prices jump are often accompanied by high levels of uncertainty, resulting in a jump of volatility. However, the jump sizes are not equal. Following Duffie *et al.* (2000), we assume the variance jump size to be exponentially distributed, i.e. $C_t \sim \exp(\mu_V)$. To allow for dependence between the jump sizes in returns and volatility, the jumps in returns are conditionally normally distributed with $Z_t | C_t \sim N(\mu_Y + \xi C_t, \sigma_Y^2)$. Due to the inclusion of the jump component in the variance process, the long-run mean changes to $\theta + \mu_V \lambda / \kappa$.

* We follow Sorensen (2002) and Richter and Sorensen (2002) and model the seasonal component using a simple sine function. An alternative approach would be to introduce monthly dummy variables in the drift as suggested by Borovkova and Geman (2007). However, as this would increase the number of parameters substantially, we prefer the more parsimonious approach based on sine functions.

23.2.2 Estimation Approach

In this section, we briefly outline the Markov chain Monte Carlo (MCMC) estimation approach we use to estimate all models. MCMC belongs to the class of Bayesian simulation-based estimation techniques. The main advantage of the MCMC methodology is the fact that it allows us to estimate the unknown model parameters and the unobservable state variables, i.e. the volatility, the jump times and the jump sizes, simultaneously in an efficient way. Jacquier *et al.* (1994) show how MCMC methods can be used to estimate the parameters and latent volatility process of a stochastic volatility model. Johannes *et al.* (1999) extend this approach for jumps in returns and, finally, Eraker *et al.* (2003) estimate the model including jumps in returns and volatility via MCMC methods.*

In order to be able to estimate the models, it is necessary to express them in discretized form. Using a simple Euler discretization, the SVCJ model is given as†

$$Y_t = Y_{t-\Delta t} + \mu(\tau)\Delta t + \sqrt{V_{t-\Delta t}}\varepsilon_t^Y + Z_t J_t \quad (23.6)$$

and

$$V_t = V_{t-\Delta t} + \kappa(\theta - V_{t-\Delta t})\Delta t + \sigma\sqrt{V_{t-\Delta t}}\varepsilon_t^V + C_t J_t. \quad (23.7)$$

The innovations ε_t^Y and ε_t^V are normal random variables, i.e. $\varepsilon_t^Y \sim N(0, \Delta t)$, and $\varepsilon_t^V \sim N(0, \Delta t)$ with correlation ρ . The jump times J_t take the value one if a jump occurs and zero if not, i.e. $J_t \sim \text{Ber}(\lambda\Delta t)$. The discretized versions of the SVJ and SV models are obtained by dropping the respective jump components. In the following, we set $\Delta t = 1$, i.e. one day.

The general idea of the MCMC methodology is to break down the high-dimensional posterior distribution into its low-dimensional complete conditionals of parameters and latent factors which can be efficiently sampled from. The posterior distribution $p(\Theta, V, J, Z, C | Y)$ provides sample information regarding the unknown quantities given the observed quantities (prices). By Bayes rule we have

$$\begin{aligned} p(\Theta, V, J, Z, C | Y) &\propto p(Y | V, J, Z, C, \Theta) \\ &p(V, J, Z, C | \Theta)p(\Theta), \end{aligned} \quad (23.8)$$

where Y is the vector of observed log prices, V, J, Z and C contain the time series of volatility, jump times and jump sizes, respectively, Θ is the vector of model parameters, $p(Y | V, J, Z, C, \Theta)$ is usually called the likelihood, $p(V, J, Z, C | \Theta)$ provides the distribution of the latent state variables, and $p(\Theta)$ the prior, reflecting the researcher's beliefs regarding the unknown parameters. To keep the influence of the priors small, we specify extremely uninformative priors.

As $p(\Theta, V, J, Z, C | Y)$ is high dimensional, it is not possible to directly sample from it. Therefore, it is necessary to simplify the problem by breaking down the posterior distribution into its complete conditional distributions which fully characterize the joint posterior. Whenever possible, we use conjugate priors which allow us to directly sample from the conditional. If this is not possible, we rely on a Metropolis algorithm. For more details on the precise specifications, see the [appendix](#).

The output of the simulation procedure is a set of G draws $\{\Theta^{(g)}, V^{(g)}, J^{(g)}, Z^{(g)}, C^{(g)}\}_{g=1:G}$ that forms a Markov chain and converges to $p(\Theta, V, J, Z, C | Y)$. Estimates of the parameters, the volatility paths, and the jump sizes are obtained by simply taking the mean of the posterior distribution. For the jump times, one additional step is required. As each draw of J is a set of Bernoulli random variables, i.e. taking on the value one or zero, the mean over all draws will provide a time series of jump probabilities. To obtain

* For an excellent introduction to MCMC estimation techniques, see Johannes and Polson (2006).

† As we work with daily data, the discretization bias is negligible.

estimates for the jump times, we follow Johannes *et al.* (1999) and identify the jump times by choosing a threshold probability, i.e. we estimate the jump times $\hat{J}$ as

$$\hat{J}_t = \begin{cases} 1, & \text{if } p(J_t = 1) > \alpha, \\ 0, & \text{if } p(J_t = 0) \leq \alpha. \end{cases} \quad (23.9)$$

The threshold α is chosen such that the number of jumps identified corresponds to the estimate of the jump intensity λ .

23.3 Data and Empirical Results

23.3.1 Data

In this paper, we employ daily spot price data for a variety of commodities traded in the US. The series are chosen to reflect the relative importance of those markets, and also to ensure the availability of as long a span of high quality data as possible. We consider two energy commodities, namely crude oil (CL) and gasoline (HU). From the metal markets, we include gold (GC) and silver (SI) in the study. Finally, we consider soybeans (S) and wheat (W) as representatives of the agricultural commodities market. All commodity data are obtained from the Commodity Research Bureau.* Additionally, we employ S&P 500 index data (obtained from Bloomberg) to enable us to put the results into the perspective of the existing literature on equity dynamics and to investigate the relationship of commodity and equity markets in Subsection 23.3.4. The data period covered is more than 25 years, spanning January 1985 to March 2010, and yielding 6290 observations per commodity and for the S&P 500.

Table 23.1 provides descriptive summary statistics and Figure 23.1 shows time series of the six commodity markets considered. Several points are worth noting. Compared with the S&P 500, the mean return is smaller for all commodities. The lowest average return is observed for the wheat market, which is barely positive, and the highest for the gold market. The standard deviation is higher than the S&P 500 for all but the gold market. The highest levels of volatility are observed for the energy commodities, with annualized values of 42.82% and 44.13%, which is more than twice the volatility of 16.00% and 18.92% observed in the gold and the equity markets. The kurtosis is, however, the highest for the S&P 500, while the lowest values are observed in the agricultural markets. However, these values are still higher than can be explained by a simple normal distribution. The smallest and largest returns are observed for the energy markets, both twice as large as for the S&P 500, which is to some extent surprising as the kurtosis levels are substantially lower. The gold and soybeans markets exhibit the smallest range between the minimum and maximum values.†

23.3.2 Univariate Analysis

Table 23.2 reports the estimation results for the SV model, i.e. the stochastic volatility model without jumps. All but the parameters affecting the mean equation and some of the correlations are statistically significant. The estimates for the speed of mean reversion, κ , are of similar size in all commodity markets and are also comparable to those of the S&P 500. Only the estimate for the silver market is somewhat higher, indicating a slightly less persistent volatility process. The long-term variance levels,

* See www.crbtrader.com.

† It is worth noting that the time series of gasoline prices exhibits a significant price spike, i.e. an upward jump with a subsequent downward jump, in the year 2005. This feature, more often observed in markets with no or limited storage possibilities (e.g. electricity or natural gas), cannot be captured very well by the models employed. We refer to Geman and Roncoroni (2006) and Nomikos and Andriosopoulos (2012) for modelling approaches including spikes. In this study, we have decided to stay within the popular class of affine jump-diffusion stochastic volatility models such that the powerful transform analysis developed by Duffie *et al.* (2000) remains applicable.




FIGURE 23.1 Price series. This figure shows the historical price series for the six commodities considered. All figures are in US dollars.

TABLE 23.1 Descriptive Statistics

	CL	HU	GC	SI	S	W	SP
Mean	0.0186	0.0189	0.0204	0.0165	0.0076	0.0008	0.0311
Std. dev.	2.6975	2.7803	1.0076	1.7621	1.5226	2.1192	1.1920
Skewness	-0.6996	-0.3748	0.1052	-1.1182	-0.6291	-0.3507	-1.3842
Kurtosis	18.0155	10.0331	11.5751	17.7078	8.1172	9.1778	32.8066
Min	-40.0011	-31.4158	-7.2327	-23.6716	-13.1820	-20.4226	-22.8997
Max	21.2765	23.0015	10.2073	13.4045	7.8667	13.1596	10.9572

Note: This table reports descriptive statistics for the daily log returns scaled by 100, i.e. Daily percentage returns. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the S&P 500. The sample period is January 2, 1985 to March 31, 2010.

θ , of all but the gold market are significantly higher than the corresponding equity estimates. The S&P 500 annualized long-term volatility is $17.58\% (= \sqrt{\theta \cdot 252})$, which is very close to the unconditional volatility, and also of similar size as in other studies.* The highest levels of annualized long-term volatility are observed for the crude oil and gasoline markets, at 36.00% and 37.61% respectively. The

* For example, Eraker *et al.* (2003) report 15.10% for the period 1985 to 1999.

TABLE 23.2 Parameter Estimates: SV

	μ	κ	θ	σ
CL	0.0575 (0.0230)	0.0204 (0.0033)	5.1432 (0.4032)	0.3676 (0.0247)
HU	0.0333 (0.0265)	0.0182 (0.0031)	5.6137 (0.4452)	0.3356 (0.0238)
GC	0.0161 (0.0082)	0.0140 (0.0025)	0.9697 (0.1243)	0.1358 (0.0085)
SI	0.0330 (0.0155)	0.0291 (0.0041)	2.7776 (0.2155)	0.3082 (0.0185)
S	0.0335 (0.0149)	0.0157 (0.0029)	2.2171 (0.2275)	0.1917 (0.0121)
W	0.0060 (0.0206)	0.0197 (0.0033)	3.8276 (0.3040)	0.2740 (0.0202)
SP	0.0334 (0.0101)	0.0177 (0.0025)	1.2261 (0.1206)	0.1597 (0.0089)
	ρ	η	ζ	
CL	-0.1266 (0.0512)			
HU	-0.0039 (0.0533)	0.0702 (0.0218)	0.0528 (0.0793)	
GC	0.3826 (0.0481)			
SI	0.3138 (0.0459)			
S	0.3206 (0.0512)	0.0460 (0.0118)	0.1363 (0.0692)	
W	0.0223 (0.0542)	0.0946 (0.0136)	0.4186 (0.0425)	
SP	-0.5678 (0.0387)			

Note: This table reports the means and standard deviations (in parentheses) of the posterior distributions for each parameter of the SV model. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the S&P 500. The sample period is January 2, 1985 to March 31, 2010.

difference across the commodity markets is considerable. The vol-of-vol parameters, σ , are highly significant, which confirms the stochastic nature of volatility in the commodity markets considered. Again, with the exception of the gold market, higher values than for the equity market are observed, implying heavier tails and greater deviance from normality. Moreover, there are substantial differences across markets—for example, the crude oil series exhibits a vol-of-vol that is twice as big as for soybeans.

Figure 23.2 plots the fitted volatility process in annualized percentages obtained from the SV model for the six commodities. It is evident that the crude oil and gasoline series are the most variable overall, while gold in particular is the least. The oil-based products showed two particular spikes in volatility during 1986 and 1990, both of which can be tied to economic events that occurred at those times. In March 1986, Saudi Arabia changed its policy and significantly expanded production to increase market share; other OPEC members followed, leading to price drops in both crude oil and gasoline. August 1990 marked the start of the first Gulf War when Iraq invaded Kuwait. Furthermore, we can observe the effect of Hurricane Katrina hitting the US Gulf coast in August 2005, leading to a short-term supply bottleneck in gasoline (although it had little noticeable effect on crude oil). Wheat was the only commodity in the sample that showed an upward trend in volatility from the late 1990s onwards. Moreover, it is also of interest to note that the Lehman default of September 2008 increased levels of volatility and uncertainty across all markets.

The estimates for the correlation of the underlying and the variance process, ρ , are most interesting, as we can observe different signs for different markets. The crude oil market shows a negative correlation, the gasoline and wheat markets (almost) zero correlations, and the gold, silver, and soybean markets significant positive correlations. For equity markets, ρ is usually found to be negative, which is confirmed by the estimate of -0.57 . The negativity of ρ in equity markets is usually explained by the leverage effect, as discussed by Black (1976). However, this argument does not apply to commodity markets. A possible explanation for the different correlations might be the fraction of hedgers and speculators in the markets. Assuming that hedging activities reduce volatility, whereas speculation activity increases volatility, this line of argument would indicate that, in the crude oil market, the fraction of hedgers over speculators increases with rising prices. In the metal and the soybean markets, the opposite would hold




FIGURE 23.2 Estimated volatility. This figure shows the estimated volatility processes for the six considered commodities under the SV model (Annualized in per cent).

true—that is, more hedging takes place for lower prices, while speculation dominates in times of higher than average prices.*

Table 23.3 provides the parameter estimates for the SVJ model, i.e. the model including Poisson price jumps. The rates of mean reversion, κ , substantially decrease in all cases and are now lower for every commodity compared with the S&P 500, implying a higher persistence of the volatility process. The long-term volatility levels (of the diffusion component) decrease, which is a technical effect, as part of the price variation is now captured by the jump component. Interestingly, however, the change for the S&P 500 is minimal compared with the commodity markets. The vol-of-vol parameters, σ , also decrease for the same reason, in that part of the excess kurtosis is now captured by the jump component. Again, the changes in the commodity markets are much bigger than for the equity case, indicating that the jump component plays an even bigger role in these markets. The correlation estimates become more pronounced compared with the SV model, i.e. the negative values decrease, whereas the positive values increase.

* We note that these hypotheses cannot be substantiated without considering the relative balance between hedgers and speculators in the respective markets. However, a detailed analysis of this aspect is beyond the scope of this paper.

TABLE 23.3 Parameter Estimates: SVJ

	μ	κ	θ	σ	ρ
CL	0.0654 (0.0221)	0.0096 (0.0019)	3.9773 (0.4431)	0.2106 (0.0159)	-0.2894 (0.0648)
HU	0.0413 (0.0274)	0.0100 (0.0027)	4.5790 (0.6264)	0.2182 (0.0222)	-0.0202 (0.0744)
GC	0.0159 (0.0081)	0.0096 (0.0019)	0.7863 (0.1201)	0.1001 (0.0059)	0.4660 (0.0591)
SI	0.0512 (0.0163)	0.0125 (0.0025)	2.0745 (0.2420)	0.1649 (0.0175)	0.5507 (0.0653)
S	0.0555 (0.0160)	0.0105 (0.0021)	1.9251 (0.2390)	0.1439 (0.0105)	0.3477 (0.0620)
W	0.0091 (0.0204)	0.0092 (0.0021)	3.2832 (0.3668)	0.1682 (0.0165)	-0.0443 (0.0727)
SP	0.0380 (0.0100)	0.0132 (0.0024)	1.2217 (0.1401)	0.1381 (0.0094)	-0.6417 (0.0369)
	λ	μJ	σJ	η	ζ
CL	0.0249 (0.0047)	-1.4218 (0.7941)	7.7162 (0.7108)		
HU	0.0241 (0.0071)	-1.1414 (0.7866)	6.5266 (0.8605)	0.0909 (0.0197)	0.0369 (0.0568)
GC	0.0531 (0.0117)	-0.0492 (0.1471)	1.8276 (0.1794)		
SI	0.0879 (0.0163)	-0.4627 (0.1901)	2.7129 (0.2476)		
S	0.0471 (0.0107)	-0.8150 (0.2658)	2.2938 (0.2473)	0.0360 (0.0177)	0.1464 (0.1169)
W	0.0225 (0.0064)	-0.7974 (0.6487)	4.7285 (0.5420)	0.0944 (0.0116)	0.4215 (0.0363)
SP	0.0071 (0.0023)	-2.5866 (0.8356)	2.8435 (0.5423)		

Note: This table reports the means and standard deviations (in parentheses) of the posterior distributions for each parameter of the SVJ model. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the s&p 500. The sample period is january 2, 1985 to march 31, 2010.

The jump intensity estimates, λ , lie between 0.022 for wheat and 0.088 for silver, providing evidence of significant differences between markets. In the oil, gasoline and wheat markets, jumps occur about six times per year. In the soybeans and gold market, we annually observe 12 to 13 jumps on average, whereas the silver price jumps about 22 times per year. These numbers are significantly larger than those observed for equity indices. In our sample period, we estimate the S&P 500 to jump only 1.8 times per year (Eraker *et al.* (2003) estimated 1.5 times per year for their sample period). The mean jump sizes, μ_j , are negative for all markets, although only significant in some instances. Taking into account the standard deviation of jump sizes, σ_j , one can observe a huge variation. A two-sigma interval covers almost as much probability mass on the positive line as on the negative. The most notable cases are the two energy markets. A plus two-sigma event in these two markets corresponds to a +14% and +12% (daily!) price jump, respectively. Analogously, a minus two-sigma event corresponds to price drops of -16% and -14%. The results for the other commodity markets are qualitatively similar, but less extreme. For the equity market, on the other hand, the mean jump size is significantly negative, and 75% of the probability mass lies on the negative line.

Lastly, Table 23.4 reports the estimates for the SVCJ models, i.e. the models including contemporaneous jumps in prices and volatility. The speed of mean reversion, κ , reverts back to values comparable to or even higher than in the SV model, indicating a lower persistence when including jumps. This makes intuitive sense as the inclusion of jumps induces the need for higher mean-reversion speeds directly after jumps (as only positive jumps are allowed). The long-term volatility levels further decrease, as part of the variation is now captured by the volatility jump component. We still observe significant differences across markets, ranging from 10% for the gold market to 27.5% for the gasoline market. The equity market lies, at 14%, closer to the lower border of this interval. The vol-of-vol parameter, σ , slightly decreases for most markets; the silver market is a notable exception, where σ increases from 0.16 to 0.21. The correlation of the diffusion components mostly decreases in absolute terms; as prices and volatility are always jumping contemporaneously, part of the correlation is now captured by the jump components. Interestingly, the jump intensities decrease significantly in some instances, such as the gold, silver, and soybean markets, whereas it remains almost unchanged in the crude oil market. Changes in the average price jump and its volatility are mostly relatively small. This is opposed to the corresponding

TABLE 23.4 Parameter Estimates: SVCJ

	μ	κ	θ	σ	ρ	λ
CL	0.0622 (0.0227)	0.0173 (0.0041)	2.4719 (0.4271)	0.2085 (0.0219)	-0.1623 (0.0623)	0.0221 (0.0043)
HU	0.0441 (0.0275)	0.0200 (0.0041)	3.0628 (0.6079)	0.2151 (0.0343)	-0.0348 (0.0721)	0.0182 (0.0063)
GC	0.0124 (0.0082)	0.0185 (0.0032)	0.4098 (0.0700)	0.0961 (0.0080)	0.3391 (0.0553)	0.0172 (0.0047)
SI	0.0283 (0.0159)	0.0374 (0.0076)	1.2995 (0.2321)	0.2110 (0.0213)	0.2583 (0.0594)	0.0305 (0.0098)
S	0.0338 (0.0157)	0.0262 (0.0046)	1.0821 (0.1656)	0.1447 (0.0158)	0.2960 (0.0657)	0.0174 (0.0056)
W	0.0164 (0.0207)	0.0240 (0.0045)	1.7065 (0.2534)	0.1373 (0.0249)	-0.0146 (0.0791)	0.0177 (0.0038)
SP	0.0421 (0.0100)	0.0225 (0.0041)	0.7874 (0.0787)	0.1264 (0.0098)	-0.5831 (0.0395)	0.0041 (0.0012)
	μJ	σJ	ξ	μV	η	ζ
CL	-1.5291 (1.1932)	7.9167 (0.7397)	-0.0170 (0.4167)	2.3176 (0.9789)		
HU	-2.0728 (1.7065)	6.7416 (0.9186)	0.1425 (0.2745)	4.0436 (1.3707)	0.0851 (0.0204)	0.0277 (0.0684)
GC	-0.3294 (0.5197)	2.7819 (0.3301)	0.5546 (0.6923)	0.5108 (0.1470)		
SI	-0.6968 (0.6375)	3.9967 (0.4951)	-0.0099 (0.3746)	1.5656 (0.6734)		
S	-0.7981 (0.8480)	2.8019 (0.3667)	-0.0612 (0.2793)	1.7129 (0.6099)	0.0456 (0.0124)	0.1455 (0.0715)
W	-1.6252 (0.9708)	4.9274 (0.5244)	0.0478 (0.2229)	3.0470 (0.7534)	0.0902 (0.0109)	0.4144 (0.0393)
SP	-4.4478 (0.9411)	2.2486 (0.5251)	0.0519 (0.2012)	2.7594 (0.7914)		

Note: This table reports the means and standard deviations (in parentheses) of the posterior distributions for each parameter of the SVCJ model. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the S&P 500. The sample period is January 2, 1985 to March 31, 2010.

result for the S&P 500, where the mean jump size almost doubles. The jump dependence parameter ξ is not estimated very precisely, a phenomenon that has already been observed by Eraker *et al.* (2003). As the values are close to zero, only modest dependence between the jump sizes is observed. The average volatility jump size lies between 0.51 for the gold, and 4.04 for the gasoline market.* Although model choice is not the main focus of this paper, in order to compare the fit of the three different models to the data, we make use of the Deviance Information Criterion (DIC) proposed by Spiegelhalter *et al.* (2002) and in particular applied to stochastic volatility models by Berg *et al.* (2004). The DIC can be regarded as a generalization of the Akaike Information Criterion (AIC), and trades off adequacy and complexity in a similar fashion.† As for the AIC, smaller DIC values indicate a better model fit.

The DIC scores are reported in Table 23.5. First, we can observe that the fit of the SV model is by far the worst for every market, providing evidence of the benefits of including a price jump component. Comparing the DIC scores of the SVJ and the SVCJ models is less conclusive as most of the scores are close to each other for a given commodity. However, the SVJ values are always lower for all commodity markets, whereas the SVCJ obtains the lowest score for the S&P 500.

A second way to compare the fit of the models to the data is by simulating price paths and comparing the average time series moments of the simulated data with those obtained from the real data. Table 23.6 reports the results of this exercise. Comparing these numbers with the descriptive statistics reported in Table 23.1, we can observe that the standard deviation is matched relatively well by the SV and SVJ models. The SV model is, however, not able to match the kurtosis levels found in the data. The SVJ model does a

* As an additional analysis, we have split the entire data set into two equal subsets and re-estimated the models for each of these. From these estimations we could observe that the most interesting parameters, such as the vol-of-vol parameters σ , the correlation ρ and the jump intensity λ do not change substantially across sub-period. Naturally, there were also some changes of the estimated parameters across the subsamples. However, it is noteworthy that these mainly affect the drift parameters, which is not surprising, and also to some extent the parameters governing the jump sizes. The latter observation needs to be interpreted with caution, as the decrease of the sample size makes the estimation of the jump component, the occurrence of which is rare by definition, quite difficult. In order to keep the presentation manageable, we do not present these results but they are available from the authors upon request.

† AIC equals DIC in the special case of flat priors.

TABLE 23.5 Model Comparison

	SV	SVJ	SVCJ
CL	27,934	26,756	26,779
HU	29,169	28,444	28,534
GC	16,096	14,634	14,871
SI	23,683	21,458	21,776
S	21,569	20,657	20,992
W	25,395	24,609	24,725
SP	17,437	16,969	16,819

Note: This table reports the DIC scores for the model specifications examined: SV, SVJ, and SVCJ. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the s&p 500. The sample period is january 2, 1985 to march 31, 2010.

TABLE 23.6 Simulated Moments

	CL	HU	GC	SI	S	W	SP
Panel A: SV							
Mean	0.0525	0.0239	0.0200	0.0330	0.0035	0.0091	0.0347
Std. dev.	2.2220	2.3116	0.9772	1.6444	1.4570	1.9325	1.0698
Skewness	-0.0299	0.0081	0.0215	0.0477	-0.0206	-0.0122	-0.0653
Kurtosis	4.0779	3.8389	3.9488	4.0809	3.8000	3.8225	3.8885
Panel B: SVJ							
Mean	0.0264	0.0173	0.0126	0.0100	0.0183	0.0104	0.0204
Std. dev.	2.3261	2.3329	0.9557	1.6184	1.4453	1.9166	1.1265
Skewness	-0.3910	-0.2842	0.0292	-0.1908	-0.2374	-0.1430	-0.4611
Kurtosis	11.6217	7.6318	5.9566	5.6459	4.7935	6.1708	6.2511
Panel C: SVCJ							
Mean	0.0094	0.0058	0.0025	0.0006	0.0062	0.0050	0.0131
Std. dev.	1.7049	1.8143	0.8876	1.1933	1.0588	1.5885	0.8706
Skewness	-0.5723	-0.3398	-0.1958	-0.2931	-0.1866	-0.3641	-0.3637
Kurtosis	16.8886	8.6686	10.9397	9.2814	6.5732	8.1980	17.0958

Note: This table reports the average first four central moments of simulated returns using the estimated parameters. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the s&p 500.

much better job and also matches the skewness closer than the SV model. In terms of skewness and kurtosis, the SVCJ model is by far the best, in many cases close to the observed levels. However, the SVCJ model does not fit the levels of standard deviations as well as the SVJ model. Overall, as the DIC values for the SVJ and SVCJ models are always close together and the analysis of simulated moments does not provide a clear conclusion as to which of the two models to prefer, we will, in the spirit of robustness, use both in the subsequent analysis. It is clear from the above analysis, however, that the jumps in returns play an important role.

23.3.3 Cross-Commodity Market Analysis

The MCMC estimation approach employed not only provides us with parameter estimates but also with estimates for the latent state variables, i.e. volatility, jump times and jump sizes. We can use this information to analyse the cross-market dependence structure of these state variables.* In a first step, we

* Alternatively, one could try to estimate a multivariate version of the model employed in order to analyse these dependencies. This approach is, however, computationally infeasible as it would involve a 14×14 covariance matrix when considering all seven markets together.

calculate the correlations of the differences in volatility, i.e. $\sqrt{V_t} - \sqrt{V_{t-1}}$, to analyse the extent to which the volatilities in the various markets move together. Table 23.7 reports these correlations, as well as return correlations to enable us to put the volatility results into perspective.

As expected, return correlations are highest between related commodities belonging to the same market segment. However, this correlation is still far from perfect. In particular, the moderate degree of correlation between crude oil and gasoline of 0.49 is interesting. The strongest correlation is observed for gold and silver with a value of 0.68. Return correlations between commodities of different segments are weak or close to zero, never exceeding 0.11.

Looking at the volatility correlations, one can identify a similar pattern. There are, however, some differences. For example, the volatility correlation between soybeans and wheat is substantially smaller than the corresponding return correlation, indicating that the prices move more closely together than the volatilities in these markets. Another interesting point is the correlation of crude oil with the agricultural commodities. The returns are mildly correlated, indicating some dependence across markets; the volatilities' correlations are, on the other hand, close to zero. Comparing the results for the SVJ and the SVCJ model, it is interesting to observe that, for some instances, the estimated correlation is almost identical (e.g. CL–HU or S–W), whereas in other instances, the result changes substantially (e.g. GC–SI). This might be a consequence of the fact that the number of jumps identified in the gold and silver markets changes substantially between the two model variants (this can be seen from the changes in the estimated value of λ).

TABLE 23.7 Correlations of Returns and Volatilities

	CL	HU	GC	SI	S	W	SP
Panel A: Returns							
CL	1.0000						
HU	0.4894	1.0000					
GC	0.0539	0.0104	1.0000				
SI	0.0404	0.0060	0.6849	1.0000			
S	0.1094	0.1098	0.0006	−0.0106	1.0000		
W	0.0969	0.0755	0.0030	−0.0157	0.3941	1.0000	
SP	0.0298	0.0459	−0.0418	−0.0630	0.0792	0.0700	1.0000
Panel B: SVJ volatilities							
CL	1.0000						
HU	0.3084	1.0000					
GC	−0.0084	0.0410	1.0000				
SI	−0.0257	0.0090	0.6375	1.0000			
S	−0.0574	−0.0157	−0.0028	0.0124	1.0000		
W	0.0368	0.1163	0.0340	0.0440	0.1146	1.0000	
SP	0.0398	0.0433	0.0470	0.0429	−0.0151	0.0418	1.0000
Panel C: SVCJ volatilities							
CL	1.0000						
HU	0.3051	1.0000					
GC	0.0392	0.0244	1.0000				
SI	0.0118	0.0274	0.4175	1.0000			
S	−0.0058	0.0041	0.0067	0.0282	1.0000		
W	−0.0003	0.0265	0.0081	0.0314	0.1045	1.0000	
SP	0.0289	0.0592	0.0403	0.0369	0.0299	0.0358	1.0000

Note: This table reports return and volatility correlations. Panel A reports the correlations of returns. Panel B displays the correlations of changes in volatility calculated from the estimated volatility process of the SVJ model. Panel C displays the correlations of changes in volatility calculated from the estimated volatility process of the SVCJ model. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat, and SP for the s&p 500. The sample period is january 2, 1985 to march 31, 2010.

Next, we analyse the simultaneous jump probabilities for each pair of commodities. To do this, we simply count the numbers of simultaneous jumps and divide it by the sample size T , i.e. we calculate the following quantity: $\sum_{t=1}^T \hat{J}_i^j \hat{J}_t^j / T$. To put these numbers into perspective, we also compute the probability of a simultaneous jump assuming independence which is given by the product of the two jump intensities. Table 23.8 provides the results.

We can observe that the simultaneous jump probabilities of commodities belonging to the same segment (i.e. the pairs CL–HU, GC–SI, and S–W) are about 4–10 times higher than one would expect under independence. For all other pairs, the probabilities are very similar, indicating that jumps of non-related commodities are independent of each other. Compared with the results for the dependence in volatility, there are some interesting differences to be observed. The simultaneous jump probability between soybeans and wheat is halved from 0.46% to 0.21% when including jumps in volatility. This is in contrast to the volatility correlation, which remains almost equal. Furthermore, one can observe a substantial decrease of simultaneous jump probabilities for most cases, whereas the correlations of volatilities remain rather constant, or even increase, when including volatility jumps.

23.3.4 Commodity and Equity Markets

We now investigate the dependencies between the individual commodity markets and the equity market, i.e. the S&P 500. It is well known that return correlations between commodities and equities are quite low, which is the reason why commodities are often considered to be a diversifier for a traditional portfolio of stocks and bonds. From Table 23.7, it can be seen that the same holds true for volatility, i.e. equity and commodity volatilities are almost uncorrelated. Consequently, commodities not only serve as a return diversifier, but as a volatility diversifier at the same time.

Table 23.8 also contains the daily simultaneous jump probabilities in equity and commodity markets as well as the corresponding probabilities assuming independence. For the SVJ model, the greatest difference is observed for the gold (0.13% vs. 0.04%) and oil (0.10% vs. 0.02%) markets, indicating that jumps in these two commodities are related to jumps in the equity market. When considering the results for the SVCJ model, one can see that the joint jump probabilities are all very small. However, compared with

TABLE 23.8 Simultaneous Jump Probabilities (%)

	CL	HU	GC	SI	S	W	SP
Panel A: SVJ							
CL	–	0.06	0.12	0.21	0.12	0.06	0.02
HU	0.40	–	0.12	0.21	0.11	0.05	0.02
GC	0.17	0.14	–	0.45	0.24	0.11	0.04
SI	0.17	0.25	2.23	–	0.43	0.20	0.07
S	0.08	0.13	0.32	0.62	–	0.11	0.04
W	0.08	0.05	0.22	0.29	0.46	–	0.02
SP	0.10	0.02	0.13	0.10	0.05	0.02	–
Panel B: SVCJ							
CL	–	0.04	0.04	0.08	0.03	0.04	0.01
HU	0.35	–	0.03	0.06	0.03	0.03	0.01
GC	0.08	0.05	–	0.06	0.03	0.03	0.01
SI	0.10	0.06	0.65	–	0.06	0.06	0.02
S	0.03	0.06	0.06	0.08	–	0.03	0.01
W	0.03	0.05	0.02	0.06	0.21	–	0.01
SP	0.03	0.03	0.05	0.06	0.06	–	0.03

Note: This table reports the simultaneous jump probabilities of returns in the lower triangular. The upper triangular reports probabilities assuming independence. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat, and SP for the s&p 500. The sample period is january 2, 1985 to march 31, 2010.

TABLE 23.9 Conditional Jump Sizes In Equity And Commodity Markets

SVJ	CL	HU	GC	SI	S	W
SP/xx	-3.60	-4.43	-4.95	-5.70	-6.51	-8.32
xx/SP	3.87	14.84	0.34	-1.11	-2.03	-2.82
SVCJ						
SP/xx	-4.39	-5.41	-5.45	-4.87	-5.05	-5.06
xx/SP	9.11	7.01	0.02	-0.54	-2.47	-5.07

Note: This table reports the average jump sizes of the s&p 500 conditional on the occurrence of a jump in a commodity market (reported in line 'SP/xx') and the average jump size in the commodity market conditional on the occurrence of a jump in the s&p 500 (reported in line 'xx/SP'). 'xx' stands for the respective commodity reported in the columns. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat, and SP for the s&p 500. The sample period is january 2, 1985 to march 31, 2010.

the previous results for the SVJ model, the differences from the case assuming independence become more pronounced, indicating that the likelihood of a simultaneous jump in volatility is higher than a simultaneous jump in prices. This is especially true for the soybeans market.

When a simultaneous jump occurs, it might be that the prices in both markets jump in the same, or alternatively, in opposite directions. Moreover, the jumps might be more or less pronounced than average jumps. To investigate this issue, we compute the average jump sizes for each commodity market conditional on the occurrence of a jump in the equity market and *vice versa*. Table 23.9 reports the results. Interestingly, the average jump size in the equity market is much more negative than the unconditional average jump size of -2.59 (SVJ model). Surprisingly, this effect is strongest for the soybean and wheat markets, i.e. a jump in the agricultural markets induces an average jump size that is two to three times bigger than normal. A consideration of the average jump sizes in the commodity markets given that the equity market jumps reveals that the energy commodities tend to jump upward if the equity market jumps downwards. This is different to their unconditional average jump size which is negative. The other commodity jumps react less to jumps in the equity market; the results for the SVCJ model are qualitatively similar but with some differences in terms of strength of the effects.

23.4 Economic Implications

23.4.1 Options Valuation

To analyse how the differences of the commodity dynamics impact upon economic quantities of interest we consider the pricing and hedging of options. One of the advantages of the continuous-time models employed in this study is that semi-closed-form options pricing formulas exist. These formulas are based on Fourier analysis, which was initially proposed by Heston (1993). Duffie *et al.* (2000) provide general solutions for all affine models and for the SVCJ model in particular.

In the following, we use the model parameters estimated in the previous section to analyse the implications for the value of European call options. As we have estimated all parameters from historical price data under the physical measure, we need to make an assumption regarding the market price of volatility and jump risk. Broadie *et al.* (2007) review the most recent literature on the sign and size of these risk premia for the S&P 500, concluding that the evidence is inconclusive; most studies report insignificant risk premia from a statistical and/or economic point of view. Due to these findings, and for simplicity, in the following we assume zero volatility and jump risk premia.*

* The results would, of course, change if this assumption were not valid, especially if the nature of the risk premia differs across commodities. This is a non-trivial question that we leave for future research.

To make the results comparable, we assume that each underlying currently trades at $S = 100$. The current variance level is assumed to be equal to its long-run average, and the risk-free rate is assumed to be zero. Using the estimated parameter values, we calculate the values of at-the-money (ATM; $X = S$), in-the-money (ITM; $X = 0.95S$), and out-of-the-money (OTM; $X = 1.05S$) call options where X denotes the strike price. Table 23.10 reports the results of this computation. The left columns report the value of the option, and the right columns express the option's value relative to the value of the corresponding ATM option.

One can observe large differences across assets. Naturally, options written on the underlyings with the highest volatility levels are most expensive. The value of the gold and S&P 500 OTM options is already close to zero (although they are only 5% out-of-the-money), whereas the values for crude oil and gasoline remain at higher levels. The differences in jump intensities, jump amplitudes and the correlations of returns with volatilities lead to some interesting differences in the relative values of ITM and OTM options across commodities.* For example, comparing the soybeans and wheat markets, one can observe that the relative value of ITM options is higher for soybeans, whereas the relative value of OTM

TABLE 23.10 Option Prices

	Absolute			Relative (%)		
	ITM	ATM	OTM	ITM	ATM	OTM
Panel A: SV						
CL	7.01	4.07	2.11	172	100	52
HU	7.14	4.26	2.32	168	100	55
GC	5.23	1.77	0.38	295	100	22
SI	6.01	3.00	1.31	200	100	44
S	5.80	2.69	1.03	215	100	38
W	6.51	3.53	1.67	185	100	47
SP	5.49	1.99	0.35	276	100	17
Panel B: SVJ						
CL	7.15	4.20	2.21	170	100	53
HU	7.17	4.28	2.33	167	100	54
GC	5.24	1.77	0.37	295	100	21
SI	6.02	2.98	1.25	202	100	42
S	5.81	2.68	0.99	217	100	37
W	6.51	3.51	1.64	185	100	47
SP	5.54	2.06	0.38	269	100	18
Panel C: SVCJ						
CL	6.65	3.56	1.66	187	100	47
HU	6.79	3.80	1.89	179	100	50
GC	5.12	1.38	0.20	372	100	14
SI	5.82	2.60	0.94	223	100	36
S	5.53	2.19	0.63	253	100	29
W	6.07	2.87	1.09	211	100	38
SP	5.38	1.71	0.19	314	100	11

Note: This table reports the option prices for a horizon of one month. The price of the underlying is normalized to 100. The strike price is set to 95 (ITM), 100 (ATM), and 105 (OTM). The left part (absolute) reports the options price, the right part (relative) reports the options price relative to the ATM option price. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the s&p 500.

* When considering implied volatilities instead of relative values, these differences would manifest themselves in the shape of the volatility smile/skew.

options is higher for wheat. Comparing the prices for the SV and SVJ with the SVCJ model, one can see that prices are lower for the latter, i.e. neglecting the possibility of volatility jumps leads to increased option prices.

To further compare the consequences of the different behavior of the commodities considered for options valuation, we analyse the values of the most popular exotic contracts, i.e. barrier options. To save space, we focus our attention on at-the-money down-and-out call and put options and vary the value of the barrier between 90% and 100% of the current spot level. To take the different levels of volatility in the markets into account, we normalize the value of the barrier option by the value of a corresponding plain-vanilla option. Consequently, all resulting values must be between zero and one.

Figure 23.3 displays the results of these computations. One can observe that, especially for the down-and-out put options, the value is quite distinct across commodities, even though we have already adjusted for the absolute volatility level by normalizing with the values of plain-vanilla options. This is due to the fact that the probability of hitting the barrier is more sensitive to the volatility and jump level than the simple option value. Inspecting the figures more closely, one can observe differences between model specifications. For example, the relative value of the down-and-out put option on wheat for a barrier of 90% is around 0.4 for the SV and SVJ models but increases to 0.5 in the SVCJ model, whereas the value of the corresponding silver option remains at 0.6 for all three models. From this result, one can deduce that the inclusion of jumps in the volatility dynamics is more relevant for wheat options. A similar observation can be made for soybeans and the energy commodities.

23.4.2 Options Hedging

Lastly, we investigate the consequences of the different stochastic behavior of commodities and different models in terms of the delta hedging of options. Consider, for example, the influence of the correlation parameter. Depending on the options position, one is long or short the underlying and long or short volatility. For example, a short put position translates into a long position in the underlying and a short position in volatility. Thus, depending on the sign of the correlation, one is either naturally diversified or not. By implementing a standard delta hedge, one neglects this relationship, which will have consequences on the hedging outcome. Similarly, the jump frequency as well as the mean and standard deviation of the jump sizes will impact upon the hedging success.

We set up our simulation analysis as follows. We consider a short position in an ATM put option with a maturity of one month. The hedging horizon τ is one week, the current variance level is assumed to be equal to the long-run mean, and interest rates are assumed to be zero. To investigate the potential hedging error, we assume that the true price dynamics are given by either the SV, the SVJ, or the SVCJ model and we calculate the corresponding options value.* We then consider a standard delta hedging strategy, i.e. we calculate the option's delta using Black's formula to set up the hedging portfolio. Next, we simulate the price and volatility dynamics until τ and calculate the new option value to obtain the pay-off of the hedging portfolio. This procedure is repeated 10,000 times to obtain a distribution of hedging errors.

Table 23.11 reports the mean, standard deviation, skewness, and kurtosis of the hedging errors. The mean hedging error is negative in all cases and larger (in absolute terms) for the energy commodities, which is as expected due to the higher volatility levels. The same observation applies to the standard deviation of the errors. When considering the skewness and kurtosis, one can observe some interesting differences across commodities and models. Considering the SV model first, we see that all distributions are skewed to the left, i.e. big negative outcomes are more likely than big positive outcomes. Interestingly, the skews for the agricultural commodities and the S&P 500 are bigger than for crude oil. Moreover, looking at the kurtosis of the hedging errors, it becomes clear that crude oil exhibits the thinnest tails, i.e. extreme negative (and positive) outcomes are less likely than for the other markets.

* As in the previous section, we assume a risk premium of zero.




FIGURE 23.3 Down-and-out options. This figure displays the value of Down-and-out barrier options as a fraction of the corresponding plain vanilla options value. The values on the x-axis correspond to the Knock-out barrier. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat and SP for the S&P 500.

The picture completely changes when introducing jumps in the price and variance processes. Whereas the mean and volatility of the hedging errors remain at similar levels, or even decrease, the skewness, and, most importantly, the kurtosis, drastically increase. This holds true for all markets, but with different rates. Under the SVCJ model, the kurtosis of crude oil rockets from 2.6 to more than 38, which is now more than twice the value observed for the S&P 500. By contrast, the rise in kurtosis for the soybeans

TABLE 23.11 Hedging Error

	CL	HU	GC	SI	S	W	SP
Panel A: SV							
Mean	-0.4441	-0.4864	-0.1964	-0.3395	-0.3018	-0.3955	-0.2319
Std	0.8305	0.8618	0.3510	0.6464	0.5272	0.7072	0.3781
Skew	-1.0725	-1.3783	-1.4260	-1.3480	-1.4964	-1.4902	-1.4691
Kurt	2.5727	4.1403	4.2224	4.1102	4.4456	4.6943	4.1992
Panel B: SVJ							
Mean	-0.4896	-0.4745	-0.1967	-0.3341	-0.3028	-0.3990	-0.2309
Std	0.8772	0.7944	0.3237	0.5446	0.4823	0.6497	0.3659
Skew	-3.4585	-2.5196	-2.1526	-2.4347	-1.8658	-2.5869	-2.0724
Kurt	22.9281	12.2694	8.5793	10.5993	5.7959	13.1202	8.3335
Panel C: SVCJ							
Mean	-0.4037	-0.4062	-0.1444	-0.2634	-0.2192	-0.2997	-0.1873
Std	0.8425	0.7275	0.2879	0.4967	0.3785	0.5319	0.3034
Skew	-4.7612	-3.4210	-3.5175	-2.9353	-2.1156	-3.7515	-2.4437
Kurt	38.8046	24.8022	23.9426	16.6159	8.5904	25.8874	15.2599

Note: This table reports the mean, standard deviation (std), skewness (skew), and kurtosis (kurt) of the hedging error when delta-hedging an atm option using black's formula. The option has a maturity of one month, and the hedging horizon is one week. CL stands for crude oil, HU for gasoline, GC for gold, SI for silver, S for soybeans, W for wheat, and SP for the S&P 500.

case is least dramatic with an increase from 4.4 to 8.6. Overall, one can make out substantial differences across commodities and models when considering the third and fourth moments of the hedging errors, which translate into a significant degree of model risk when delta hedging in these markets.

23.5 Conclusions

This paper has examined the stochastic behavior of the prices and volatilities of a sample of six of the most important commodity markets. Using a Bayesian Markov chain Monte Carlo estimation approach, three separate stochastic volatility-type models are estimated and compared for each commodity price series, and for the S&P 500 by way of comparison. Fairly intuitively, correlations between the returns are high for pairs of commodities from the same sub-class but almost zero across sub-classes. The same pattern holds for the relationships between the simultaneous jump probabilities: within market segments, jumps often occur in tandem, whereas they are essentially independent across segments.

We are able to demonstrate that not only are return correlations between commodities and the stock index low, as is well documented, but the correlations between commodity and stock volatilities are also low. This is an important result since it shows that commodities may be an even more useful portfolio constituent than previously thought as they can act as a volatility diversifier as well, which is important for options portfolios or any portfolio of securities with embedded contingent claims. The paper examines whether jumps occur simultaneously across commodities and equities; there is some evidence that jumps in the two asset classes do occur together, notably for gold and oil, although the results vary somewhat between models.

Finally, we test the economic impact of employing one stochastic volatility model rather than another by considering the pricing of European calls and barrier options written on each commodity, and by evaluating the effectiveness of delta hedging. Again, we find substantial differences both across commodities and between models, indicating the heterogeneous nature of this asset class and the importance of judiciously choosing the most appropriate specification for each individual series.

Acknowledgements

We thank two anonymous referees of Quantitative Finance and discussants and participants at the German Finance Association Meeting (DGF), Regensburg, 2011, the Meeting of the Swiss Society for Financial Research (SGF), Zurich, 2012 and the European Meeting of the Financial Management Association (FMA), Istanbul, 2012, for helpful comments.

References

- Aiube, F., Baidya, T. and Tito, E., Analysis of commodity prices with the particle filter. *Energy Econ.*, 2008, **30**, 597–605.
- Alizadeh, A., Nomikos, N. and Pouliasis, P., A Markov regime switching approach for hedging energy commodities. *J. Bank. Finance*, 2008, **32**, 1970–1983.
- Andersen, T., Benzoni, L. and Lund, J., An empirical investigation of continuous-time equity return models. *J. Finance*, 2002, **62**, 1239–1284.
- Asgharian, H. and Bengtsson, C., Jump spillover in international equity markets. *J. Financ. Econometr.*, 2006, **4**, 167–203.
- Bakshi, G., Cao, G. and Chen, C., Empirical performance of alternative option pricing models. *J. Finance*, 1997, **52**, 2003–2049.
- Ball, C. and Torous, W., On jumps in common stock prices and their impact on call option pricing. *J. Finance*, 1985, **40**, 155–173.
- Bates, D., Jumps and stochastic volatility: exchange rate processes implicit in Deutschmark options. *Rev. Financ. Stud.*, 1996, **9**, 69–108.
- Bates, D., Post-'87 crash fears in S&P 500 futures options. *J. Econometr.*, 2000, **94**, 181–238.
- Berg, A., Meyer, R. and Yu, J., Deviance information criterion for comparing stochastic volatility models. *J. Business Econ. Statist.*, 2004, **22**, 107–120.
- Black, F., Studies of stock price volatility changes. In Proceedings from the American Statistical Association, Business and Economic Statistics Section, pp. 177–181, 1976 (American Statistical Association).
- Borovkova, S. and Geman, H., Seasonal and stochastic effects in commodity forward curves. *Rev. Deriv. Res.*, 2007, **9**(2), 167–186.
- Brennan, M. and Schwartz, E., Evaluating natural resource investments. *J. Business*, 1985, **58**, 135–157.
- Broadie, M., Chernov, M. and Johannes, M., Model specification and risk premia: evidence from futures options. *J. Finance*, 2007, **62**, 1453–1490.
- Casassus, J. and Collin-Dufresne, P., Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance*, 2005, **60**, 2283–2331.
- Chernov, M., Ghysels, E., Gallant, A. and Tauchen, G., Alternative models for stock price dynamics. *J. Econometr.*, 2003, **116**, 225–257.
- Du, X., Yu, C. and Hayes, D., Speculation and volatility spillover in the crude oil and agricultural commodity markets: a Bayesian analysis. *Energy Econ.*, 2011, **33**, 497–503.
- Duffie, D., Pan, J. and Singleton, K., Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 2000, **68**(6), 1343–1376.
- Eraker, B., Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *J. Finance*, 2004, **59**, 1367–1403.
- Eraker, B., Johannes, M. and Polson, N., The impact of jumps in volatility and returns. *J. Finance*, 2003, **58**, 1269–1300.
- Erb, C. and Harvey, C., The strategic and tactical value of commodity futures. *Financ. Anal. J.*, 2006, **62**, 69–97.
- Geman, H. and Nguyen, V.-N., Soybean inventory and forward curve dynamics. *Mgmt Sci.*, 2005, **51**, 1076–1091.
- Geman, H. and Roncoroni, A., Understanding the fine structure of electricity prices. *J. Business*, 2006, **79**, 1225–1261.
- Gibson, R. and Schwartz, E., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.

- Gorton, G. and Rouwenhorst, G., Facts and fantasies about commodity futures. *Financ. Anal. J.*, 2006, **62**(2), 47–68.
- Haigh, M. and Holt, M., Crack spread hedging: accounting for time-varying volatility spillovers in the energy futures markets. *J. Appl. Econometr.*, 2002, **17**(3), 269–289.
- Heston, S., A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.*, 1993, **6**, 327–343.
- Jacquier, E., Polson, N. and Rossi, P., Bayesian analysis of stochastic volatility models. *J. Business Econ. Statist.*, 1994, **12**, 371–389.
- Johannes, M., Kumar, R. and Polson, N., State dependent jump models: How do US equity indices jump. Working Paper, 1999.
- Johannes, M. and Polson, N., MCMC methods for financial econometrics. In *Handbook of Financial Econometrics*, edited by Y. Ait-Sahalia and L. Hansen, 2006 (Elsevier: Amsterdam).
- Kat, H. and Oomen, R., What every investor should know about commodities, Part I: univariate return analysis. *J. Invest. Mgmt*, 2007a, **5**, 4–28.
- Kat, H. and Oomen, R., What every investor should know about commodities, Part II: multivariate return analysis. *J. Invest. Mgmt*, 2007b, **5**, 40–64.
- Koop, G., *Bayesian Econometrics*, 2003 (Wiley: New York).
- Larsson, K. and Nossman, M., Jumps and stochastic volatility in oil prices: time series evidence. *Energy Econ.*, 2011, **33**, 504–514.
- Liu, P. and Tang, K., The stochastic behavior of commodity prices with heteroskedasticity in the convenience yield. *J. Empir. Finance*, 2011, **18**, 211–224.
- Manoliu, M. and Tompaidis, S., Energy futures prices: term structure models with Kalman filter estimation. *Appl. Math. Finance*, 2002, **9**, 21–43.
- Merton, R., Option pricing when underlying stock returns are discontinuous. *J. Financ. Econ.*, 1976, **3**, 125–144.
- Ng, V. and Pirrong, C., Price dynamics in refined petroleum spot and futures markets. *J. Empir. Finance*, 1996, **2**, 359–388.
- Nomikos, N. and Andriosopoulos, K., Modelling energy spot prices: empirical evidence from nymex. *Energy Econ.*, 2012, **34**, 1153–1169.
- Pindyck, R., Volatility and commodity price dynamics. *J. Fut. Mkts*, 2004, **24**, 1029–1047.
- Richter, M. and Sorensen, C., Stochastic volatility and seasonality in commodity futures and options: the case of soybeans. Working Paper, 2002.
- Sadorsky, P., Modeling and forecasting petroleum futures volatility. *Energy Econ.*, 2006, **28**, 467–488.
- Samuelson, P., Proof that properly anticipated prices fluctuate randomly. *Ind. Mgmt Rev.*, 1965, **6**, 41–49.
- Schwartz, E., The stochastic behavior of commodity prices: implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.
- Schwartz, E. and Smith, J., Short-term variations and long-term dynamics in commodity prices. *Mgmt Sci.*, 2000, **46**, 893–911.
- Scott, L., Option pricing when the variance changes randomly: theory, estimation, and an application. *J. Financ. Quantit. Anal.*, 1987, **22**, 419–438.
- Serletis, A., A cointegration analysis of petroleum futures prices. *Energy Econ.*, 1994, **16**, 93–97.
- Sorensen, C., Modeling seasonality in agricultural commodity futures. *J. Fut. Mkts*, 2002, **22**(5), 393–426.
- Spiegelhalter, D., Best, N., Carlin, B. and van der Linder, A., Bayesian measures of model complexity and fitting. *J. R. Statist. Soc. B*, 2002, **64**, 583–639.
- Tang, K., Time-varying long run mean of commodity prices and the modelling of futures term structure. *Quant. Finance*, 2012, **12**, 781–790.
- Trolle, A. and Schwartz, E., Unspanned stochastic volatility and the pricing of commodity derivatives. *Rev. Financ. Stud.*, 2009, **22**, 4423–4461.
- Wang, T., Wu, J. and Yang, J., Realized volatility and correlation in energy futures markets. *J. Fut. Mkts*, 2008, **28**, 993–1011.

Appendix A: MCMC estimation details

This appendix provides more detailed information on the MCMC estimation procedure. For a general introduction to MCMC techniques, see, for example, Koop (2003). The Gibbs sampling technique allows one to draw each parameter and state variable of the joint posterior sequentially. For many parameters, conjugate priors can be used to derive the conditional posterior distribution, which is a standard distribution and therefore easily sampled from. Details are given below.

- μ we use a standard normal distribution as the prior
- κ, θ we use a truncated normal distribution bounded at zero and hyperparameters of 0 and 1 as priors for each parameter
- σ, ρ as the posteriors for σ and ρ are not known, we follow the re-parameterization suggested by Jacquier *et al.* (1994) and use the inverse Gamma distribution with hyperparameters 1 and 200 as priors
- η we use an exponential distribution with a hyperparameter of 0.05 as a prior, yielding a truncated normal distribution as a posterior
- ζ the posterior distribution for ζ is non-standard, and we apply an independence Metropolis algorithm, drawing from the uniform distribution on the unit interval
- λ we use a Beta distribution with hyperparameters 2 and 40
- μ_V we use a Gamma distribution with hyperparameters 1 and 1
- μ_J we use a Normal distribution with hyperparameters 0 and 100 as priors
- σ_J^2 we use an Inverse Gamma distribution with hyperparameters 5 and 20 as priors
- ξ we use a standard normal distribution as a prior

The posteriors of J_t , Z_t and C_t are non-standard but are provided in the appendix of Eraker *et al.* (2003). The conditional posterior distribution of the variance path V is also non-standard and given by

$$p(V | Y, J, Z, C, \Theta) \propto \prod_{t=1}^T p(V_t | V_{t-1}, V_{t+1}, Y, J, Z, C, \Theta), \quad (\text{A.1})$$

where the posterior function for each V_t is given as

$$p(V_t | V_{t-1}, V_{t+1}, Y, J, Z, C, \Theta) \propto \frac{1}{V_t} e^{-(1/2)(\omega_1 + \omega_2 + \omega_3)}, \quad (\text{A.2})$$

with

$$\omega_1 = \frac{(Y_{t+1} - \mu(\tau) - J_{t+1}Z_{t+1})^2}{V_t}, \quad (\text{A.3})$$

$$\omega_2 = \frac{(V_t - \theta\kappa - (1-\kappa)V_{t-1} - \sigma\rho(Y_t - \mu(\tau) - J_tZ_t) - J_tC_t)^2}{\sigma^2 V_{t-1} (1-\rho^2)}, \quad (\text{A.4})$$

$$\omega_3 = \frac{(V_{t+1} - \theta\kappa - (1-\kappa)V_t - \sigma\rho(Y_{t+1} - \mu(\tau) - J_{t+1}Z_{t+1}) - J_{t+1}C_{t+1})^2}{\sigma^2 V_t (1-\rho^2)}. \quad (\text{A.5})$$

We use a random walk Metropolis algorithm. The volatility of the error term in this procedure is calibrated such that the acceptance probability is within the range 30–50%. See Koop (2003) for details of calibrating the Random Walk Metropolis algorithm. Each parameter (and state variable) is sampled 100,000 times (i.e. $G = 100,000$) and we discard the first 30,000 ‘burn-in’ draws as is standard practice with MCMC estimations.

A Hybrid Commodity and Interest Rate Market Model

24.1	Introduction	523
24.2	The Commodity LIBOR Market Model	525
	The Interest Rate Market • The Commodity Market	
24.3	Calibration With Time-Dependent Volatilities	527
24.4	Futures/Forward Relation and Convexity Correction	530
24.5	Merging Interest Rate and Commodity Calibrations	531
24.6	Real Data Example	534
24.7	Pricing Spread Options	541
24.8	Conclusion	546
	Acknowledgements	546
	References	546
	Appendix A: Volatility integrals	547

Kay F. Pilz

Erik Schlögl

A joint model of commodity price and interest rate risk is constructed analogously to the multi-currency LIBOR Market Model (LMM). Going beyond a simple ‘re-interpretation’ of the multi-currency LMM, issues arising in the application of the model to actual commodity market data are specifically addressed. Firstly, liquid market prices are only available for options on commodity futures, rather than forwards, thus the difference between forward and futures prices must be explicitly taken into account in the calibration. Secondly, we construct a procedure to achieve a consistent fit of the model to market data for interest options, commodity options and historically estimated correlations between interest rates and commodity prices. We illustrate the model by an application to real market data and derive pricing formulas for commodity spread options.

Keywords: Applied mathematical finance; Arbitrage pricing; Calibration of deterministic volatility; Commodity markets; Computational finance; Derivative pricing models; Interest rate modeling; Interest rate derivatives

JEL Classification: C6, C63, G1, G13

24.1 Introduction

Traditionally, market risks in different asset classes, such as FX, fixed income or commodities, have been modelled separately, although modelling approaches incorporating several sources of market risk are becoming increasingly popular. The aim of this paper is to contribute to the more integrated approach by applying the multi-currency LIBOR Market Model (LMM) as presented by Schlögl (2002b) jointly to interest rates and commodities. The domestic fixed income market will be interpreted in the usual way

for the LMM, i.e. with a given bond market paying in a certain currency (say USD), whereas the foreign market will be associated with a commodity market (e.g. crude oil), with the physical commodity as its currency, and the ‘convenience yield’ as its rate of interest. The contribution of this paper goes beyond a simple ‘re-interpretation’ of the multi-currency LMM in that issues arising in the application of the model to actual commodity market data are specifically addressed. Firstly, liquid market prices are only available for options on commodity futures, rather than forwards, thus the difference between forward and futures prices must be explicitly taken into account in the calibration. Secondly, we construct a procedure to achieve a consistent fit of the model to market data for interest rate options, commodity options and historically estimated correlations between interest rates and commodity prices. This is achieved in a two-stage process, calibrating first to the interest rate and commodity markets separately, followed by an orthonormal transformation of the commodity volatility vectors to ‘rotate’ the commodity volatilities relative to the interest rate volatilities in such a manner as to achieve the desired correlations between the two markets. Thirdly, we illustrate the use of the model on actual market data and demonstrate a way of efficiently calculating commodity spread options.

For the market data illustration, we chose WTI Crude Oil as the commodity and US dollars as the ‘domestic’ currency. Thus the ‘exchange rate’ is the WTI Crude Oil price, i.e. the WTI Crude Oil futures quotes—when converted to forward prices using an appropriate convexity correction—can be interpreted as forward exchange rates between the US dollar economy and an economy where value is measured in terms of units of WTI Crude Oil (and where convenience yields play the role of interest rates). In this example, the model is calibrated to USD at-the-money caplets and swaptions, as well as WTI Crude Oil futures and ATM options on futures.

The model presented here builds on the prior literature on the LMM, starting from the seminal work of Miltersen *et al.* (1997), Brace *et al.* (1997) and Jamshidian (1997). The model construction follows Musiela and Rutkowski (1997) and the calibration method builds on the approach first proposed by Pedersen (1998) for the basic LMM.

The earliest work combining commodity and interest rate risk (based on dynamics of the continuously compounded short rate, without reference to a full model calibrated to an initial term structure) goes back to Schwartz (1982). Subsequently, models for commodity price dynamics incorporating stochastic convenience yields have been constructed by a number of authors, for example Gibson and Schwartz (1990), Cortazar and Schwartz (1994), Schwartz (1997), Miltersen and Schwartz (1998), and Miltersen (2003). These models are typically constructed on the basis of a Heath *et al.* (1992) term structure model with generalized (possibly multi-factor) Ornstein/Uhlenbeck dynamics, i.e. continuously compounded convenience yields (and possibly interest rates rates) are normally distributed. While in theory these models allow some freedom to calibrate to market data, the published work does not contain much evidence as to whether an effective calibration to available commodity and interest rate options—even at the money—can be achieved. This type of fit to observed market prices, in particular to at-the-money options, is a strength of the LIBOR Market Models that we exploit in the present paper.

The paper is organized as follows. The basic notation, the results of the single- and multi-currency LMM and their interpretation in the context of commodities are presented in [Section 24.2](#). In [Section 24.3](#) the calibration of the commodity part of the Commodity LMM to plain vanilla options is discussed. In all major commodity markets, only options with the same maturity as the underlying forward are traded liquidly. Hence, no volatility term structure for a particular forward commodity price can be extracted only from the options written on this forward. To overcome this problem, option prices for all maturities are used to calibrate simultaneously a volatility term structure, which is desired to be as time-homogeneous and smooth as possible. In keeping with the assumptions of the LMM, we thus focus on the volatility term structure only, i.e. with time dependence but no strike dependence. In [Section 24.4](#) the relationship between futures and forwards in the model is presented, which permits calibration of the model to futures as well as forwards. The calibration of the interest rate part of the hybrid Commodity LMM will not be discussed in detail in this paper, because this problem has already been addressed by many authors and most methods should be compatible with our model. However,

in [Section 24.5](#) we discuss how both separately calibrated parts—the interest rate and the commodity part—of the model can be merged in order to have one underlying d -dimensional Brownian motion for the joint model and still match the market prices used for calibration of the particular parts. [Section 24.6](#) illustrates the application of the Commodity LMM to real market data. Finally, [Section 24.7](#) addresses the pricing of commodity spread options in the model.

24.2 The Commodity LIBOR Market Model

24.2.1 The Interest Rate Market

The construction of the LMM for the domestic interest rate market follows the presentation of Musiela and Rutkowski (1997). Assume a given probability space $(\Omega, \{\mathcal{F}_t\}_{t \in [0, T^*]}, \mathbb{P})$, where the underlying filtration $\{\mathbb{F}_t\}_{t \in [0, T^*]}$ coincides with the $\mathcal{P}$ -augmentation of the natural filtration of a d -dimensional standard Brownian motion W . Let T^* be a fixed time horizon, then a family of bond prices is any family of strictly positive real-valued adapted processes $B(t, T)$ for $t \in [0, T]$, with $B(T, T) = 1$ for every $T \in [0, T^*]$. The bond price $B(t, T)$ for $T \in [t, T^*]$ is the amount that has to be invested at time t to receive one unit of the domestic currency at time T .

Corresponding to assumptions (BP.1) and (BP.2) of Musiela and Rutkowski (1997), we have that the bond price $B(t, T)$ for $t \leq T \leq T^*$ is a strictly positive special semi-martingale (see Musiela and Rutkowski (1997, p. 263) for the definition) and that the forward process,

$$F_B(t, T, T^*) := \frac{B(t, T)}{B(t, T^*)}, \quad (24.1)$$

follows a martingale under the T^* -forward measure $\mathbb{P}$. Equivalently, the bond price process satisfies

$$B(t, T) = \mathbb{E}_{\mathbb{P}} \left[\frac{B(t, T^*)}{B(T, T^*)} \middle| \mathcal{F}_t \right], \quad (24.2)$$

for all $t \in [0, T]$.

We present some conclusions from these assumptions, which will be of use later and can be found in Musiela and Rutkowski (1997). Firstly, there exists a $\mathbb{R}^d$ -valued process $\gamma(t, T, T^*)$, such that the forward process has the representation

$$dF_B(t, T, T^*) = F_B(t, T, T^*) \gamma(t, T, T^*) \cdot dW_{T^*}(t). \quad (24.3)$$

By Itô's formula and Girsanov's theorem it follows that, for any given $S, T \in [0, T^*]$, the dynamics of $F_B(t, S, T)$ for $t \in [0, S \wedge T]$ can be written as

$$dF_B(t, S, T) = F_B(t, S, T) \gamma(t, S, T) \cdot dW_T(t), \quad (24.4)$$

with

$$\begin{aligned} \gamma(t, S, T) &= \gamma(t, S, T^*) - \gamma(t, T, T^*), \\ dW_T(t) &= dW_{T^*}(t) - \gamma(t, T, T^*) dt. \end{aligned} \quad (24.5)$$

Hence, $F_B(t, S, T)$ is an exponential (local) martingale under the $\mathbb{P}_T$ measure, and since both $\mathbb{P}$ and $\mathbb{P}_T$ are absolutely continuous, the Radon–Nikodým derivative is given by the Doléans exponential

$$\frac{d\mathbb{P}_T}{d\mathbb{P}} = \varepsilon_T \left(\int_0^\cdot \gamma(u, T, T^*) \cdot dW_{T^*}(u) \right), \quad \mathbb{P}_{T^*}\text{-a.s.}$$

The measure $\mathbb{P}_T$ is usually denoted as the T -forward measure, and $\mathbb{P}_{T^*} = \mathbb{P}$ additionally as the terminal measure.

For practical purposes it is often convenient to use a discrete-tenor version of the LMM.* Therefore, we assume that the time horizon $T^* = N\delta$ is a multiple of a fixed period δ . Then, the LIBOR forward rate $L(t, T)$ as seen at time t , for an investment of one currency unit at time T with pay-off $B(t, T)/B(t, T + \delta)$ at $T + \delta$, can be written as

$$dL(t, T) = L(t, T)\lambda(t, T) \cdot dW_T(t), \quad (24.6)$$

for $t \in [0, T]$. The volatility function of LIBOR relates to the volatility function of the forward process by

$$\gamma(t, T, T + \delta) = \frac{\delta L(t, T)}{1 + \delta L(t, T)} \lambda(t, T). \quad (24.7)$$

To shorten notation, define $T_i = i\delta$ (for $0 \leq i \leq N$).

24.2.2 The Commodity Market

The approach incorporating a commodity as followed in this section parallels the multi-currency LMM introduced by Schlögl (2002b). The commodity market can naturally be considered as a ‘foreign interest rate market’, where the currency is the physical commodity itself, e.g. as in our example in [Section 24.6](#), the value of any asset is measured in units (barrels) of crude oil instead of dollars. The bond prices $C(t, T)$ quote (as seen at time t) the amount of the commodity that has to be invested at time t to physically receive one unit of the commodity at time T . Corresponding LIBORs can be interpreted in the same way, i.e. the yield of $C(t, T)$ is the convenience yield (adjusted for storage costs, if applicable). Although these rates have a natural interpretation, they are usually not traded (liquidly) for most of the commodities. We will discuss below how this model can still be calibrated to the instruments commonly traded in the commodities market.

Some assumptions are required with regard to the commodity market. Firstly, all assumptions made in constructing the LMM on the domestic interest rate market in the previous section are assumed to be true for the commodity market as well. This particularly includes the process $C(t, T)$ for $0 \leq t \leq T \leq T^*$ to be adapted to the filtration $\{\mathcal{F}_t\}_{t \in [0, T^*]}$. We will denote the corresponding T_i -forward measures, Brownian motions and volatility functions by $\mathbb{P}_{T_i}$, $\tilde{W}_{T_i}$ and $\tilde{\gamma}(t, T_i, T_{i+1})$, respectively.

Secondly, we postulate the existence of a spot exchange rate process $X(t)$, which is a positive special semi-martingale under $\mathbb{P}_{T^*}$. $X(t)$ is the spot price of the commodity at time t .

A commodity bond $C(t, T)$ converted by the spot exchange rate is then a traded asset in the domestic market (denominated in domestic currency), and hence its forward value,

$$X(t, T_i) = \frac{C(t, T_i)X(t)}{B(t, T_i)} \quad (0 \leq i \leq N), \quad (24.8)$$

is a martingale under $\mathbb{P}_{T_i}$ and is called the T_i -forward exchange rate, and in the present context this is the forward price of the commodity. If the dynamics of the forward exchanges rates are written in terms of

$$dX(t, T_i) = X(t, T_i)\sigma_X(t, T_i) \cdot dW_{T_i}(t) \quad (0 \leq i \leq N), \quad (24.9)$$

in which the volatility functions are not necessarily deterministic, then it is shown by Schlögl (2002b) that, under no-arbitrage restrictions, the volatility functions must satisfy the relation

$$\sigma_X(t, T_{i-1}) = \tilde{\gamma}(t, T_{i-1}, T_i) - \gamma(t, T_{i-1}, T_i) + \sigma_X(t, T_i) \quad (1 \leq i \leq N). \quad (24.10)$$

* Where an extension to continuous tenor is required, it is typically more practical to start with a discrete-tenor version of the LMM and extend it using ‘interpolation by daycount fractions’ as in Schlögl (2002a) in order to avoid the infinite-dimensional state variables resulting from the continuous-tenor extensions proposed by Brace *et al.* (1997) and Musiela and Rutkowski (1997).

This generally leaves two ways to link the interest rate and commodity markets in order to have an arbitrage-free multi-currency LMM. The first would be to calibrate the single-currency LMM models separately, i.e. determine the volatility functions $\gamma(t, T_p, T_{i+1})$ and $\tilde{\gamma}(t, T_i, T_{i+1})$ (or $\lambda(t, T_i)$ and $\tilde{\lambda}(t, T_i)$ equivalently) for $i = 1, \dots, N-1$, and then calibrate the forward exchange rate volatility $\sigma_X(\cdot, T_i)$ for one arbitrary forward time T_i . Using (24.10), all other forward exchange rate volatilities $\sigma_X(\cdot, T_j)$ for $1 \leq j \neq i \leq N$ can be derived. This approach seems especially appropriate when linking two (real) interest rate markets (like USD and AUD), because LIBORs (or similar) and currency forwards are liquidly traded for all major currencies.

In order to link an interest rate market with a commodity market, the second approach seems to be more appropriate. The fixed income market volatilities $\gamma(t, T_p, T_{i+1})$ (or, equivalently, $\lambda(t, T_i)$) are calibrated as for a single-currency LMM—in our example in Section 24.6, this will be the USD market calibrated to at-the-money caplets and swaptions. Then, for all forward times $T_0, \dots, T_N$ volatility functions $\sigma_X(\cdot, T_i)$ are calibrated to the commodity options market (in Section 24.6, we specifically consider options on WTI Crude Oil futures), i.e. the forward commodity prices are assumed to follow log-normal dynamics under the appropriate probability measure, rather than the simple-compounded convenience yields.* The volatility functions $\tilde{\gamma}(t, T_i, T_{i+1})$ and $\tilde{\lambda}(t, T_i)$ can now be derived from relations (24.10) and (24.7), respectively.

Remark 1: If a deterministic volatility function σ_X is chosen for the commodity forward process, it will be shown in Section 24.4 that the corresponding futures process is not log-normally distributed, apart from the futures price at maturity. Nevertheless, the difference between forward and futures prices can be expressed in terms of volatilities for the commodity forwards and forward interest rates.

24.3 Calibration with Time-Dependent Volatilities

Since the Commodity LMM is based on commodity forwards, we have to calibrate to forward implied volatilities or plain vanilla option prices written on forwards. However, commodities futures rather than forwards are most liquidly traded (consider, for example, the WTI Crude Oil futures in the market data example in Section 24.6) and thus forward prices have to be deduced from the futures. As we are specifically concerned with integrating commodity and interest rate risk, it is not adequate to equate forward prices with futures prices, as is still common among practitioners. The following Section 24.4 describes how to approximate the difference between futures and forwards as well as the implied forward volatilities, in order to apply the calibration methods proposed in the present section.

We first state some notational conventions and modify slightly the notation in Section 24.2, adapted from the prior literature. Since in our setup the forward exchange rate $X(t, T)$ is a commodity forward, we will write $F(t, T)$ instead, where t is termed the process time (aka calendar time) and T the forward time. Accordingly, the volatility of the forwards will be denoted by $\sigma(t, T)$, instead of $\sigma_X(t, T)$, and is the instantaneous volatility at time t of the forward with forward time T .

We assume we have forward processes $F(\cdot, T_1), \dots, F(\cdot, T_N)$ with expiries $T_1, \dots, T_N$ and we further think of T_0 as ‘now’. Times to maturity for an arbitrary calendar time $t \geq 0$ are given by $x_i = T_i - t$ for $i = 0, 1, \dots, N$. Finally, market prices for call options on $F(\cdot, T_i)$ with pay-off $[F(T_p, T_i) - K]^+$ are assumed to be available and denoted by $C_{i,\text{mkt}}^+$ for $i = 1, \dots, N$.

For calibration we further need to determine vectors of calendar times $t_c = (0, t_1, \dots, t_{n_c})$ and times to maturity $x_f = (x_0, x_1, \dots, x_{n_f})$, which define a grid for $V = (v_{i,j})_{1 \leq i \leq n_c, 1 \leq j \leq n_f}$, the matrix of piecewise constant instantaneous volatilities. The entry $v_{i,j}$ represents the volatility corresponding to forward $F(t, T)$ with $t_{i-1} \leq t < t_i$ and $x_{j-1} \leq T - t < x_j$. The number of forward times n_f in the volatility matrix need not coincide with the number of traded forwards N , and especially in regions of large forward times a rougher spacing can be chosen for x_f , since volatilities tend to flatten out with increasing forward time.

* Assuming the convenience yields to be log-normal also does not appear reasonable in light of the fact that, empirically, convenience yields can become negative.

Let $v(t, T)$ denote the (annualized) implied volatility as seen at time t for forward time T , then the relation between time-dependent instantaneous and implied volatilities is given by

$$\frac{1}{T-t} \int_t^T \sigma^2(s, T) ds = v^2(t, T). \quad (24.11)$$

Identity (24.11) is utilized to compute a (model) option price C_i^{mod} (for $i = 1, \dots, N$) from the volatility matrix V with Black's formula (Black, 1976). Note that the assumption of deterministic volatilities for commodity forwards made in the previous section permits us to apply Black's formula in a consistent manner. In Appendix A we show how to compute the integral when volatilities are piecewise constant. Now, the first calibration criterion measures the *quality of fit* by considering the squared differences between market and model call prices. In order to weight the various fit criteria we assign different weights η to each of them,

$$q = \eta_q \sum_{j=1}^N (C_j^{\text{mkt}} - C_j^{\text{mod}})^2. \quad (24.12)$$

As already mentioned in the introduction, in commodity markets, typically only options with maturities coinciding with the maturity of the underlying forward exist. In order to construct a full volatility term structure for a particular forward, it is necessary to use information gained from other options on forwards expiring before the particular one. Following Pedersen (1998), to this end several heuristic concepts can be employed, which more or less rely on market practice and experience. We will refer to these as *smoothness criteria*. The first is time-homogeneity, which assumes that the (unknown) volatility at future calendar time $t = T_2 - T_1$ of a forward $F(\cdot, T_2)$ having then time to maturity T_1 and time to maturity T_2 now, will be 'similar' to the (known) volatility the forward $F(\cdot, T_1)$ with time to maturity T_1 has now. In other words, under strict time-homogeneity, $\sigma(t, T)$ can be expressed as a function of time to maturity $T - t$ only. This is usually too restrictive to be able to fit the model to the market in general, so instead of enforcing strict time-homogeneity, we penalize departures from time-homogeneity in the first smoothness criterion. In terms of the volatility matrix V , this means that volatilities with different calendar times but the same time to maturity should not differ too much, so the penalty function is

$$s_1 = \eta_1 \sum_{j=1}^{n_t} \sum_{i=1}^{n_t-1} (v_{i+1,j} - v_{i,j})^2. \quad (24.13)$$

Another typical characteristic is the so-called Samuelson effect, which states that forward volatilities tend to decrease with increasing time to maturity. In order to be consistent with this behaviour, the volatility matrix has to be monotonically decreasing in forward time, and violations of this property are penalized by

$$s_2 = \eta_2 \sum_{i=1}^{n_t} \sum_{j=1}^{n_t-1} (\max\{v_{i,j+1} - v_{i,j}, 0\})^2. \quad (24.14)$$

Finally, we enable a criterion that imposes the volatility term structure to be smooth in time to maturity for each fixed calendar time. Assigning a large weight to this criterion would force the volatility to be flat in the forward time direction, which is usually not desirable, however with a small weight this criterion contributes to a smoother volatility surface,

$$s_3 = \eta_3 \sum_{i=1}^{n_t} \sum_{j=1}^{n_t-1} (v_{i,j+1} - v_{i,j})^2. \quad (24.15)$$

To minimize the objective value $q + s_1 + s_2 + s_3$, a least-squares optimization method is used, one that allows for non-linear dependencies between model parameters and objective function values.

Remark 2: The quality-of-fit criterion q and the smoothness criteria s_1, s_2, s_3 are constructed such that the absolute squared differences between model and market option prices and between neighbouring volatilities, respectively, are minimized. Alternatively, relative deviations could be used for minimization, but we did not experience substantial differences in the quality of fit or the smoothness of volatilities when doing so.

Remark 3: Regarding the relationship between calendar times and forward times of the volatility matrix and the times to maturity of the available options on forwards, two important points should be noted. Firstly, to be able to price options on all of the forwards, the maturity of the longest available forward has to be smaller than or equal to the latest calendar time and the longest time to maturity, $T_N \leq \min\{t_{n_t}, x_{n_t}\}$. Here, we assume that all options mature at the same time as their underlying forwards. Secondly, in the context of the integral in (24.11), volatilities $\sigma(t, T)$ at any calendar time t with forward time $T > T_N$ or, equivalently, $t + x > T_N$, have no impact on the price of any plain vanilla option used for calibration. In terms of the piecewise constant volatility matrix, this means that volatilities $v_{i,j}$ with $t_{i-1} + x_{j-1} > T_N$ have no contribution to the quality-of-fit criterion and are therefore determined only by smoothness criteria.

We want to conclude this section by briefly describing the concept of volatility factor decomposition and factor reduction, which can either be included in the calibration process, or applied to the final volatility matrix after calibration. As will be demonstrated in the following Sections 24.4 and 24.5, volatility factor decomposition has to be included in the calibration process when calibrating to futures and options on futures. The method has to be applied separately for every calendar time, which is therefore fixed to arbitrary t_i in the following, and let v_i denote the i th row (corresponding to calendar time t_i) of V , written as a column vector. Together with the exogenously given correlation matrix C , which is assumed to be constant over calendar time, the covariance matrix is calculated by

$$\Sigma = (v_i v_i^\top) \odot C, \quad (24.16)$$

where $\top$ means transposition and $\odot$ multiplication by components (Hadamard product). Σ is symmetric and positive definite, hence can be decomposed into $\Sigma = R D^{1/2} (R D^{1/2})^\top$. The columns of $R = (r_{j,k})_{1 \leq j, k \leq n_t}$ consists of orthonormal eigenvectors and the diagonal matrix $D = (\xi_{j,k})_{1 \leq j, k \leq n_t}$ of the eigenvalues of Σ . This representation allows us to reduce the number of stochastic factors, because usually the first d eigenvalues (when ordered decreasingly in D and the columns of R accordingly) account for more than around 95% of the overall variance, with d substantially smaller than n_t . Hence, R and D can be reduced to matrices $R \in \mathbb{R}^{n_t \times d}$ and $D \in \mathbb{R}^{d \times d}$ by retaining only the first d columns in R and the upper $d \times d$ sub-matrix in D , respectively. Instead of v_i we now use the matrix $U = (u_{j,k})_{1 \leq j \leq n_t, 1 \leq k \leq d}$ (the calendar time index i is omitted for notational convenience) with entries $u_{j,k} = r_{j,k} \sqrt{\xi_k}$, which relates to v_i by

$$v_{i,j}^2 = \sum_{k=1}^d u_{j,k}^2 = \sum_{k=1}^d r_{j,k}^2 \xi_k \quad (j = 1, \dots, n_t). \quad (24.17)$$

The computation of the variance in (24.11), as described in Appendix A, allows also for representations of the volatilities in decomposed form.

Remark 4: The correlation matrix C in (24.16) is the correlation matrix of the forward returns, not of the futures returns.* However, for the convexity adjustment suggested in Section 24.4, both correlation matrices coincide.

Remark 5: When used as in (24.16), the calculation of a correlation matrix from a historical time series should be consistent with the forward time concept of the volatility matrix. That means that if the

* Strictly speaking, it is the matrix of quadratic covariation of the forward processes.

volatility matrix is given for calendar times t and absolute forward times T , the correlations should be calculated from instruments with the same absolute maturity times. Correspondingly, if the volatility matrix is given for calendar times t and times to maturity x , the correlations should be calculated from instruments with the same times to maturity. Typically, futures are quoted in the market for absolute maturity times and forwards (if they are quoted at all) have constant times to maturity. Although, in practice, both methods typically do not exhibit substantial differences, an interpolation could be performed before estimating the correlations in order to switch from absolute maturity times to times to maturity or *vice versa*. For a discussion of different correlation concepts and their effect on calibration and pricing in the context of the LMM, see Choy *et al.* (2004).

24.4 Futures/Forward Relation and Convexity Correction

The calibration method in Section 24.3 is applicable only when forwards and options on forwards are available. This section presents the approximate conversion of futures prices to forward prices for all relevant data for calibration, in order to apply the methods of the previous section when only futures and options on futures are available (such as in the case of WTI Crude Oil as considered in the market data example in Section 24.6).

We introduce the notation $G(t, T)$ for a futures price at time t with maturity T , and, as before, $F(t, T)$ will be the corresponding forward price. From no-arbitrage theory we know that $F(T, T) = G(T, T)$ and that prices of plain vanilla options on forwards and futures must coincide, whenever the maturities of the option, forward and futures are the same. This allows us to use the call prices of options on futures for calibration of forwards, and we only have to ensure that the (virtual) forwards have the same maturities as the futures. Due to Equation (24.9) the forward $F(\cdot, T_i)$ is an exponential martingale under the T_i -forward measure, and since it has deterministic volatility, it is log-normally distributed under $\mathbb{P}_{T_i}$. A change of measure from the T_i -forward to the spot risk-neutral measure relates the Brownian motions by

$$dW_{T_i}(t) = \eta(t, T_i)dt + dW_{\mathbb{Q}}(t) \quad (1 \leq i \leq N),$$

and $\eta(\cdot, T_i)$ relates to the volatility of the process $F_B(t, T_i, T_i + \delta)$ by

$$\begin{aligned} \eta(t, T_{i+1}) - \eta(t, T_i) &= \gamma(t, T_i, T_{i+1}) \\ &= \frac{\delta L(t, T_i)}{1 + \delta L(t, T_i)} \lambda(t, T_i) \quad (1 \leq i \leq N-1), \end{aligned} \quad (24.18)$$

with $\eta(t, T_1) = 0$ if we are using ‘interpolation by daycount fractions’ to extend the model to continuous tenor as in Schlögl (2002a). Hence, for $1 \leq i \leq N$,

$$\begin{aligned} F(T_i, T_i) &= F(t, T_i) \exp \left\{ \int_t^{T_i} \sigma(u, T_i) \cdot dW_{T_i}(u) - \frac{1}{2} \int_t^{T_i} \|\sigma(u, T_i)\|^2 du \right\} \\ &= F(t, T_i) \exp \left\{ \int_t^{T_i} \sigma(u, T_i) \cdot dW_{\mathbb{Q}}(u) - \frac{1}{2} \int_t^{T_i} \|\sigma(u, T_i)\|^2 du \right\} \\ &\quad \times \exp \left\{ \int_t^{T_i} \sigma(u, T_i)^\top \eta(u, T_i) du \right\}, \end{aligned}$$

where $\|\cdot\|$ denotes the inner product norm. Furthermore, futures follow the general relation

$$G(t, T) = \mathbb{E}_{\mathbb{Q}}[S(T) | \mathcal{F}_t] \quad (24.19)$$

(see, e.g. Cox *et al.* (1981) and Miltersen and Schwartz (1998)), where $S(t)$ is the spot price that satisfies by no-arbitrage constraints $S(t) = F(t, t) = G(t, t)$ for all t . Putting these relations together, the futures in (24.19) can be expressed as

$$\begin{aligned} G(t, T_i) &= \mathbb{E}_{\mathbb{Q}}[F(T_i, T_i) | \mathcal{F}_t] \\ &= F(t, T_i) \mathbb{E}_{\mathbb{Q}} \left[\exp \left\{ \int_t^{T_i} \sigma(u, T_i) \cdot dW_{\mathbb{Q}}(u) - \frac{1}{2} \int_t^{T_i} \|\sigma(u, T_i)\|^2 du \right\} \right. \\ &\quad \left. \times \exp \left\{ \int_t^{T_i} \sigma(u, T_i)^\top \eta(u, T_i) du \right\} \middle| \mathcal{F}_t \right], \end{aligned} \quad (24.20)$$

and is obviously not log-normally distributed, since $\eta(\cdot, T_i)$ depends on the forward interest rates. The difficulty here is to calculate an expectation of a random process value under a measure for which the process is not a martingale. Similar techniques as utilized for convexity correction in interest rate theory can be applied in order to obtain an expression under the expectation operator that is a log-normally distributed random variable with respect to the $\mathbb{Q}$ -measure.*

A simple and widely used way is to make the second term in the expectation deterministic by ‘freezing’ the level dependence of $\eta(\cdot, T_i)$ with respect to the currently observed forward curve, i.e. defining $\bar{\eta}(\cdot, T_i)$ by

$$\bar{\eta}(t, T_{i+1}) - \bar{\eta}(t, T_i) = \frac{\delta L(0, T_i)}{1 + \delta L(0, T_i)} \lambda(t, T_i). \quad (24.21)$$

Then all volatility terms in (24.20) are deterministic and we have

$$G(t, T_i) = F(t, T_i) \exp \left\{ \int_t^{T_i} \sigma(u, T_i)^\top \bar{\eta}(u, T_i) du \right\} \quad (1 \leq i \leq N). \quad (24.22)$$

In Appendix A we show how to compute the integral in (24.22) when the volatility functions $\sigma(t, T)$ and $\bar{\eta}(t, T)$ are piecewise constant.

Remark 6: Based on the convexity correction (24.22), the identity of the correlation matrices for forward and futures returns follows immediately.

24.5 Merging Interest Rate and Commodity Calibrations

So far the interest rate market and the commodity market have been considered separately, with the exception of the convexity correction in (24.22) that involved the forward interest rate volatility. In this section we focus on linking both volatility matrices for building a joint Commodity LMM. The linkage is controlled by the correlation matrix between the interest rate and commodity forwards. The whole correlation matrix, i.e. for any pair of forwards or forward rates within the set of commodity forwards and forward interest rates, is assumed to be constant over time. Further, we assume that the calendar time discretization, for which the matrices of the piecewise constant volatility for commodities and interest rates have been calibrated, coincide. The forward time discretization may differ.

The method to be described has to be applied separately for every calendar time, which is therefore fixed to some t_i in the following. The same techniques as applied for factor reduction in the context of [Section 24.3](#) will be employed here. We add to the notation of the aggregate volatility matrix V from [Section 24.3](#) the subscript ‘C’ for ‘commodity’ in order to distinguish it from its interest rate equivalent, which is denoted by ‘I’. Further, we decompose the volatility matrices V_C (obtained from the calibration

* See, for example, Pelsser (2000, chapter 11) for details on convexity correction.

in [Section 24.3](#)) and V_I (obtained from the LMM calibration, for example by the method described by Pedersen (1998)) to matrices

$$U_C = (u_{j,k}^C)_{1 \leq j \leq n_t, 1 \leq k \leq d_C}, \quad U_I = (u_{j,k}^I)_{1 \leq j \leq m_t, 1 \leq k \leq d_I},$$

where m_t is the number of interest forward rates and d_C and d_I are the number of factors with significant eigenvalues for commodity forwards and interest forward rates, respectively. The columns of each matrix correspond to the stochastic factors and the rows to different forward times. Moreover, each column vector is orthogonal to the other column vectors within the same matrix.

We start to link both calibrations by choosing a parameter $d \geq d_C, d_I$ that determines the number of relevant factors for both commodities and interest rates. d will also be the dimension of the vector of Brownian motions in the joint Commodity LMM. If d is chosen to be greater than d_C or d_I , respectively, the corresponding matrices $U_C \in \mathbb{R}^{n_t \times d_C}$ and $U_I \in \mathbb{R}^{m_t \times d_I}$ have to be enlarged to matrices $U_C \in \mathbb{R}^{n_t \times d}$ and $U_I \in \mathbb{R}^{m_t \times d}$, simply by adding zero columns at the end. The forward interest rates now depend only on factors that relate to non-zero columns, but it is different for commodity forwards. Our fitting procedure, presented below, subsequently modifies U_C and therefore determines which factors will have an impact solely on the commodity forwards, which will have no impact on the commodity forwards and which will contribute to both commodity forwards and forward interest rates, and therefore to the cross-correlation.

The separate calibrations for commodities and interest rates are merged by matching the model intrinsic cross-covariance matrix $U_C U_I^\top$ with the cross-covariance matrix calculated from the exogenously given cross-correlation matrix C_{CI} ,

$$\Sigma_{CI}^{\text{target}} = \sqrt{\text{diag}\{v_{Ci} v_{Ci}^\top\} \text{diag}\{v_{Ii} v_{Ii}^\top\}^\top} \odot C_{CI},$$

where diag (applied to a matrix) returns the diagonal of the matrix as column vector and the square root has to be applied component-wise. As in [Section 24.3](#), v_{Ci} and v_{Ii} denote the i th row (corresponding to calendar time t_i) of V_C and V_I , respectively, written as column vectors. In order to achieve $U_C U_I^\top \approx \Sigma_{CI}^{\text{target}}$ we exploit the property of multivariate normal distributed random variables to be invariant to orthonormal rotations. That means we have to find a matrix Q satisfying $QQ^\top = I_d$ (where I_d denotes the $d \times d$ -identity matrix) that minimizes

$$r_1 = \zeta_1 \| \Sigma_{CI}^{\text{target}} - U_C Q U_I^\top \|,$$

with respect to some matrix norm, e.g. the Frobenius norm. The weight factor ζ_1 may be necessary when further constraints have to be controlled by the loss function as well, as will be discussed below. Alternatively, it is also possible to define the cross-correlation matrix to be the target matrix and minimize

$$r_1 = \zeta_1 \| C_{CI} - (U_C Q U_I^\top) \odot (\text{diag}\{v_{Ci} v_{Ci}^\top\} \text{diag}\{v_{Ii} v_{Ii}^\top\}^\top)^{-1/2} \|,$$

subject to the same orthonormality constraint for Q . The exponent $-1/2$ has to be applied component-wise.

As with calibration, the optimization for obtaining a Q can be performed by any non-linear optimization procedure that allows for non-linear constraints, or by non-linear least-squares algorithms such as Levenberg/Marquardt, in which the distance of $QQ^\top$ to the identity matrix I_d is again controlled by a matrix norm,

$$r_2 = \zeta_2 \| QQ^\top - I_d \|,$$

ensuring that Q is as close as possible to a valid orthonormal rotation preserving the original interest rate and commodity calibrations.

Remark 7: The cross-correlations between commodity forwards and interest rate forwards are much lower than the correlations within the asset classes, and estimation from historical data appears to be

much more volatile for the cross-correlations than for the correlations within the asset classes. For example, the structure of the cross-correlation matrix between WTI Crude Oil forwards and USD interest rate forwards in Figure 24.7 of Section 24.6 can hardly be explained by obvious rationales. Therefore, in practice, one might wish to specify a flat cross-correlation founded on particular market views. This can be realized in a straightforward manner in our approach, but the example in Section 24.6 below shows that it is also possible to adequately match a more complicated cross-correlation structure. In the end, there is a trade-off between the quality of the cross-correlation fit and the number of stochastic factors in the model. In view of computational efficiency and model parsimony a limited number of factors is desirable and perhaps more important than exactly fitting a given cross-correlation structure.

Remark 8: In our approach the basis transformation Q applies to the commodity volatility matrix U_C . Alternatively, one could choose the interest rate volatility matrix U_I for transformation. This would not change criteria r_1 and r_2 , because Q can also be interpreted as Q^T with inverse Q , and the transformed model covariance matrix would read as $U_C(U_I Q)^T$.

The determination of the basis transformation matrices concludes the calibration of the model, if the market instruments are forwards and options on forwards. However, if the calibration is carried out for futures and options on futures, the model option prices would change with any non-trivial basis transformation, since the forward prices are modified by means of the convexity correction (24.22).

This implies that the whole calibration process has to be iterated by, first, refitting the volatility matrix V_C , where in each step of the optimization the corresponding basis transformation matrix Q has to be multiplied by the volatility decomposition U_C (for each calendar time) before the model call prices and the value of the loss function are computed. Secondly, for the refitted V_C and U_C , new basis transformations have to be determined in order to still match the cross-correlations.

Both steps, the fitting of V_C and Q , have to be iterated until a sufficient smoothness and quality of fit subject to the market option prices and the exogenously given cross-correlation is reached. For the real data example of the following section, this occurs already after three iterations in the case with six stochastic factors.

For clarity we summarize the steps of the whole calibration process to futures and options on futures.

(I) Preliminary calculations applied to the LMM calibration outcome.

1. Computations for each calendar time t_i ($1 \leq i \leq n_i$).
 - a. Computation of the covariance matrix Σ_{Ii} as in (24.16).
 - b. Decomposition of Σ_{Ii} into U_{Ii} using PCA in the way described at the end of Section 24.3.

(II) Iteration until a sufficient quality of fit based on the criteria on matching market option prices, smoothness and cross-correlation is reached.

1. Calibration of V_C .
Minimization of the penalty function given in (d):
 - a. Computations for each calendar time t_i ($1 \leq i \leq n_C$).
 - i. Computation of the covariance matrix Σ_{Ci} by (24.16).
 - ii. Decomposition of Σ_{Ci} into U_{Ci} using PCA as described at the end of Section 24.3.
 - iii. Multiplication by the basis transform Q_i resulting from the previous iteration, $U_{Ci}Q_i$ (Q_i is the identity matrix in the very first iteration).
 - b. Computation of forward prices from market futures prices using (24.22).
 - c. Computation of model prices for options on forwards using (24.17) and the Black formula.
 - d. Calculation of the loss value $q + s_1 + s_2 + s_3$ as defined in Section 24.3.
2. Fit of the basis transformations.
Fitting $\{Q_i\}_{1 \leq i \leq n_C}$ subject to the penalty function $r_1 + r_2$ as described in this section.

Remark 9: The number of factors d_C can affect the number of iterations required for the calibration process to converge satisfactorily. Equations (24.21) and (24.22) demonstrate that if the cross-correlation

is zero, futures and forward prices coincide. On the other hand, the first fit of the volatility matrices U_C is made without any consideration of cross-correlation. Hence (unintentionally) assigning in the computation of U_C the most contributing eigenvalues and eigenvectors of V_C to those stochastic factors of the d -dimensional Brownian motion, which also represent strong V_I contributions, would generate a rather high cross-correlation in the first calibration step at the end of Step II.1 above. In this case the forward curve would tend to depart from the futures curve. However, if the exogenous cross-correlation has a low level, the basis transformation would modify the volatility matrix U_C in a way that not only reduces the high model cross-correlation generated in the first calibration step, but also returns the forward curve in the direction of the futures curve by reducing the magnitude of the convexity correction. An iteration of the procedure (as described above) will typically (in all real data scenarios that we have considered) force the objective variables V_C and Q to a balanced state with adequate calibration results. As demonstrated in the following section, the convergence is much slower if d_C is large relative to d_I , since a large d_C allows for more de-correlation than a small value.

Finally, in the last iteration of the procedure described above we end up with decomposed volatility matrices U_C and transformation matrices Q for each calendar time, and assembling these resulting n_c -many matrices $U_C Q$ and U_I into three-dimensional arrays $\Lambda_C \in \mathbb{R}^{n_c \times n_t \times d}$ and $\Lambda_I \in \mathbb{R}^{n_c \times m_t \times d}$ the calibration of the hybrid Commodity LMM is finished.* Note that Λ_I basically consists of piecewise constant forward interest rate volatilities $\lambda(t, T)$ as occurring in (24.6), obtained from the separate LMM calibration, at most modified by some interpolation on these calibrated volatilities in order to obtain Λ_I . Λ_C is the result of the calibration method as described in the previous sections, together with the transformation presented in this section in order to match cross-correlations.

This allows us to write the dynamics of the hybrid Commodity LMM as follows. Let W be a d -dimensional Brownian motion and denote by $\lambda_{i,j,\cdot}^I$ and $\lambda_{i,j,\cdot}^C$ the d -dimensional vectors in Λ_I and Λ_C of volatilities for calendar times $t \in [t_{i-1}, t_i]$ and times to maturity $x \in [x_{j-1}, x_j]$. Then, the dynamics of the forward interest rates $L(t, T)$ as given in (24.6) can be written as

$$dL(t, T) = L(t, T) \lambda_{i,j,\cdot}^I \cdot dW_T(t),$$

and the dynamics of the commodity forwards $F(t, T)$ as given in (24.9) as

$$dF(t, T) = F(t, T) \lambda_{i,j,\cdot}^C \cdot dW_T(t)$$

for all maturity times T satisfying $T - t \in [x_{j-1}, x_j]$ and all calendar times $t_{i-1} \leq t < t_i$ (for some $1 \leq i \leq n_c$ and $1 \leq j \leq m_t$ or $1 \leq j \leq n_t$, respectively).

24.6 Real Data Example

We demonstrate the performance and applicability of the Commodity LIBOR Market Model by calibrating it to real data.† We have chosen May 5, 2008 for calibration and WTI Crude Oil as commodity, hence the US Dollar (USD) forward rate as interest rate. As can be seen in Figure 24.1, the nearest WTI Crude Oil future price of 119.97 USD was not too far from its peak in July 2008. The futures curve is in ‘backwardation’ and covers a rather large range of about 10 USD within the first five years of maturity. The 3-month forward rates show a less extreme pattern than the commodity futures. An application to 2009 data produced results comparable to those presented in this example regarding the fit to commodity and fixed income market data, but at the cost of a slightly rougher forward interest rate volatility surface.

* Since we had to equalize the calendar time discretization for commodities and interest rates in order to be able to fit a transformation matrix Q , the first dimension of Λ_I is of size n_c . The choice to adapt the forward interest rate discretization of calendar time to the commodity discretization is arbitrary and any other discretization could be used.

† The data were taken from the SuperDerivatives platforms SD-IR and SD-CM.

The calibration of the (classical) LMM was done as proposed by Pedersen (1998). Figure 24.2 shows the resulting volatility surface and the correlation matrix as used for calibration, which was estimated historically from the time series of forward rates covering the 3 months before May 5, 2008. Caplets, caps and swaptions were used for calibration and the fit to market prices is quite good, as Figure 24.3 demonstrates. Note that we have to employ the forward interest rate curve up to 6 years forward in time, in order to calculate the convexity correction for a commodity volatility surface with 3 years calendar and 3 years forward time. The calibration of the WTI Crude Oil futures was achieved by the method described in Section 24.3. The market instruments are futures and options on futures traded on the New York Mercantile Exchange. Figure 24.4 shows the calibrated volatility surface and the historically estimated correlation matrix using 3 months of futures prices before the calibration date. Calendar and forward times go out to 3 years, and although on the exchange futures with expiries in every month are traded, we chose the calendar and forward time vectors to be unequally spaced (while still calibrating to all traded instruments), with 1 month difference up to 1 year, 2.4 months difference between 1 and 2 years




FIGURE 24.1 The commodity and interest market for calibration date May 5, 2008. Left: The WTI Crude Oil nearest futures between 2005 and the end of 2008. The circle indicates the calibration date. Middle: The futures curve as seen at calibration date with maturities up to five years. Right: The 3-month USD forward rates for reset dates (expiries) between 3 months and 4 years and 9 months.




FIGURE 24.2 Left: The calibrated volatility matrix. Right: The historically estimated correlation matrix as used for calibration.




FIGURE 24.3 Market prices versus model prices. The market prices are represented by crosses connected by solid lines, the model prices by circles connected by dashed lines. Left: Caplets. Middle: Caps. Right: Swaptions with 1 year, 2 year and 3 year tenors (from bottom to top).

and 6 months difference between 2 and 3 years. This speeds up the calibration without losing much structure in the volatility surface, since the market views futures with long maturity as having almost flat volatility. For weighting in the calibration objective function we have chosen $\eta_q = 1$ (fit to call prices), $\eta_1 = 0.1$ (time homogeneity, i.e. smoothness in the calendar time direction), $\eta_3 = 0.01$ (smoothness in the forward time direction) and $\eta_2 = 0.1$ (decreasing monotonicity in the forward time direction). As initial guess we have interpolated the implied volatilities to the forward time grid for the first calendar time t_1 and then used constant extrapolation to all other calendar times.

The model fit to the commodity call prices is very good, as illustrated in the right graph of Figure 24.5. The left graph shows the difference between futures and forwards as calculated by the convexity correction.

In Figure 24.6, particular slices of the forward volatility structure are examined. The left graph shows the term structure of volatilities for certain calendar times, whereas the right graph illustrates the evolution of volatilities over the lifetime of different forwards.

Finally, we link both separately calibrated volatility matrices to one set of stochastic factors. Table 24.1 shows how much of the overall variance, i.e. of the sum of variances over all factors, can be explained by the leading factors, when the factors are sorted according to decreasing contribution to total variance of the commodity forwards. The first two factors already account for more than 99% of the overall variance, hence a reasonable choice would be $d_c = 2$. Since Pedersen's calibration method for the LMM part of the model already includes a spectral decomposition of the interest forward rates covariance matrix, it is not meaningful to do it again here, because the number of factors with a reasonably large contribution should be the same as the number of factors chosen for the LMM calibration. If it is necessary to interpolate the forward interest rate volatility matrix in order to match the calendar times of the commodity volatility matrix, the forward rate covariance matrix will change and, hence, eigenvalue decompositions of the calendar time adjusted covariance matrices yield different results than eigenvalue decompositions of the original covariance matrices as used for calibration. However, these differences should not be substantial as long as the calibrated volatility matrix is sufficiently smooth in calendar time. For the LMM calibration we have chosen the number of factors to be $d_1 = 4$, which again covers about 99% of the overall variance.

We now have to determine the parameters for the final cross-correlation fit. First, we choose the total number of factors in the joint model to be $d = 6$ and allow the commodity volatility matrix to disperse over all six factors, whereas the interest rate volatilities should remain on the first four factors, since we apply the basis transform to the commodity volatilities (see also Remark 8). Figure 24.7 demonstrates the result of the transformation. The left graph shows two surfaces, the exogenously given target cross-correlation matrix (as calculated from the historical time series covering 3 months before




FIGURE 24.4 Left: The calibrated commodity volatility surface. Right: The historically estimated correlation matrix as used for calibration.




FIGURE 24.5 Left: The differences between futures (dashed red line) and forwards (solid blue line) as calculated from the convexity correction. Right: Fit of model prices (red circles connected by dashed line) to market prices (blue crosses connected by solid line).




FIGURE 24.6 Left: Forward volatility structures at different calendar times (solid black = 0 (calibration date), dotted-dashed blue = 6 month, dashed red = 1 year, dotted green = 2 years). Right: Evolution of forward volatilities over the lifetime (solid black = forward with 3 years maturity, dashed-dotted blue = 2 years, dashed red = 1 year, dotted green = 6 month).

TABLE 24.1 Results of the eigenvalue decomposition of the commodity forward covariance matrix for the first calendar time t_1 . The maximum number of factors coincides with the number of forward times greater than zero. The second row shows, for the i th factor, the percentage of overall variance that can be generated by the first i factors

	No. of factors							
	1	2	3	4	5	6	...	19
Percentage of overall variance	98.138	99.835	99.957	99.979	99.988	99.992	...	100

the calibration time) and the transformed correlation matrix for the particular calendar time that fitted worst with respect to the Frobenius norm. The right graph illustrates the absolute differences between the target cross-correlation and the fitted cross-correlation matrix of the left graph. For optimization we applied a non-linear Levenberg/Marquardt algorithm with scale parameters $\zeta_1 = 1$ (quality of fit) and $\zeta_2 = 10$ (orthonormality of the transformation matrix). Without the application of a basis transformation, i.e. by optimizing only subject to the quality of fit and smoothness criteria, the cross-correlation ranges between 0.5 and 1 for all commodity and interest rate forwards over all calendar times.

By construction, the aggregated commodity volatility matrix, as shown in the left graph of Figure 24.4, does not change. How the decomposed factors change is illustrated in Figure 24.8. The left graph shows only the first four factors of the original calibration without considering cross-correlations, since the others are very close to zero. The right graph demonstrates how the basis transformation spreads the contribution of the relevant factors of the original calibration to all available factors in order to match the cross-correlation coefficients.




FIGURE 24.7 The cross-correlation fit with $d_C = d = 6$. Left: The target cross-correlation matrix (coloured) as estimated from historical futures returns and the cross-correlation matrix for the calendar time that fitted worst (grey). Right: Differences between target and fitted cross-correlations.




FIGURE 24.8 The factorized commodity volatilities of the first calendar time t_1 . Left: The first four factors from the initial calibration, without considering cross relations. The fifth and sixth factors are not shown since they are almost zero. Right: All six factors after applying the basis transformation in order to match the cross-correlations. The fifth and sixth factors are represented by those two lines that have the largest value at forward time closest to zero.

The quality of the cross-correlation fit can be substantially improved by increasing d_C and hence d . Allowing for six more factors, i.e. setting $d_C = d = 12$, results in cross-correlation fits as shown in Figure 24.9.

The choice of $d_C - d_I = 8$ independent factors for the commodity forwards allows the fitted basis transformation matrices in the very first iteration to have a strong de-correlation effect. This is especially true compared with the high cross-correlation (between 0.5 and 1.0) obtained from fitting V_C without any consideration of cross-correlation in step II.1 of the very first iteration. Table 24.2 demonstrates that, for $d_C = 12$, more iterations are required to find a balance for the fits of V_C and the cross-correlation matrix.

Figure 24.10 shows the forward prices and the fit to the market call prices for both calibrations. Whereas the forwards almost coincide, the fit to market call prices seems to be worse in the case of $d_C = 12$, especially for options with long maturities. The left graph of Figure 24.11, which is a detail of the right graph of Figure 24.10, confirms this and explains the larger loss value ℓ_1 . The smoothness of the volatility matrix V_C is lower for $d_C = 12$, as is demonstrated by the right graph of Figure 24.11. This further contributes to a higher ℓ_1 .




FIGURE 24.9 The cross-correlation fit with $d_C = d = 12$. Left: The target cross-correlation matrix (coloured) as estimated from historical futures returns and the cross-correlation matrix for the calendar time that fitted worst (grey). Right: Differences between target and fitted cross-correlations.

TABLE 24.2 Calibration with $d_C = 6$ and three iterations compared with $d_C = 12$ and 12 iterations. The value ℓ_1 denotes the final loss function value of the V_C optimization in step II.1. The value ℓ_{2i} denotes the Frobenius norm of the difference between the model and target cross-correlation matrix for calendar time t_i as induced by the fitted Q_i in step II.2

		Iteration					
		1	2	3	4	5	6
$d_C = 6$	ℓ_1	0.0261	0.0265	0.0263	–	–	–
	$\min\{\ell_{2i}\}$	0.2373	0.2409	0.2414	–	–	–
	$\max\{\ell_{2i}\}$	0.3544	0.3922	0.3972	–	–	–
$d_C = 12$	ℓ_1	0.0261	0.9852	0.9550	0.7856	0.7828	0.6816
	$\min\{\ell_{2i}\}$	0.2004	0.2005	0.2005	0.2000	0.2000	0.2000
	$\max\{\ell_{2i}\}$	0.2551	0.2636	0.2543	0.2879	0.2868	0.2539
		Iteration					
		7	8	9	10	11	12
	ℓ_1	0.6367	0.6301	0.6327	0.5090	0.4146	0.3992
	$\min\{\ell_{2i}\}$	0.2001	0.2000	0.2000	0.2001	0.2012	0.2012
	$\max\{\ell_{2i}\}$	0.2538	0.2537	0.2537	0.2538	0.2539	0.2537




FIGURE 24.10 Comparison of the calibration results for $d_c = 6$ with three iterations and $d_c = 12$ with 12 iterations. Left: The market future prices (red dashed line) and the forward prices as calibrated with $d_c = 6$ (blue solid line) and $d_c = 12$ (green dotted-dashed line). Right: The market call prices (black circles) and the model call prices for $d_c = 6$ (solid blue line) and $d_c = 12$ (green dotted-dashed line).




FIGURE 24.11 Comparison of the calibration results for $d_c = 6$ with three iterations and $d_c = 12$ with 12 iterations. Left: A detail of the right graph of Figure 24.10. The market call prices (black circles) and the model call prices for $d_c = 6$ (solid blue line) and $d_c = 12$ (green dotted-dashed line). Right: Forward volatilities as calibrated with $d_c = 12$ for different calendar times (solid black = 0 (calibration date), dotted-dashed blue = 6 months, dashed red = 1 year, dotted green = 2 years). See the left graph of Figure 24.6 for the corresponding volatilities calibrated with parameter $d_c = 6$.

24.7 Pricing Spread Options

This section examines an approximation approach in order to price spread options on forwards and futures following Kirk (1995). Spread options can be used to hedge the risk of roll over costs, i.e. the price difference between two futures or forwards with different maturities. Moreover, spread options depend on the correlation and the difference between the volatilities of the involved futures or forwards, and

hence can be employed to reduce the risk generated by these parameters. We will first consider forwards and then sketch an analogous derivation for futures.

Let T_i, T_j denote two forward maturities with $T_i < T_j$ and we set the option expiration time to T_i . According to market practice, the spread is defined by $S_{ij}(t) = F(t, T_j) - F(t, T_i)$ and the spread call option value is given by

$$\begin{aligned} C_{\text{Spread}}^{\text{Fwd}}(0, T_i, T_j, K) &= \mathbb{E}_{\mathbb{Q}}[D(T_i)(S_{ij}(T_i) - K)^+ | \mathcal{F}_0] \\ &= B(0, T_i) \mathbb{E}_{T_i}[(F(T_i, T_j) - F(T_i, T_i) - K)^+ | \mathcal{F}_0]. \end{aligned} \quad (24.23)$$

From Section 24.2 we know that

$$dF(t, T_i) = F(t, T_i) \sigma(t, T_i) \cdot dW_{T_i}(t), \quad (24.24)$$

$$\begin{aligned} dF(t, T_j) &= F(t, T_j) \sigma(t, T_j)^\top \left(\sum_{\ell=i}^{j-1} \frac{\delta L(t, T_\ell)}{1 + \delta L(t, T_\ell)} \lambda(t, T_\ell) \right) dt \\ &\quad + F(t, T_j) \sigma(t, T_j) \cdot dW_{T_i}(t). \end{aligned} \quad (24.25)$$

By ‘freezing’ the level dependence of $L(\cdot, T_\ell)$ with respect to the currently observed forward curve as described in Section 24.4, both processes become geometric Brownian motions,

$$dF(t, T_i) = F(t, T_i) \sigma(t, T_i) \cdot dW_{T_i}(t), \quad (24.26)$$

$$\begin{aligned} dF(t, T_j) &\approx F(t, T_j) \sigma(t, T_j)^\top \Gamma_{i,j-1}(t) dt \\ &\quad + F(t, T_j) \sigma(t, T_j) \cdot dW_{T_i}(t), \end{aligned} \quad (24.27)$$

with

$$\Gamma_{i,j-1}(t) = \sum_{\ell=i}^{j-1} \frac{\delta L(0, T_\ell)}{1 + \delta L(0, T_\ell)} \lambda(t, T_\ell). \quad (24.28)$$

The idea of Kirk (1995) is to assume that, instead of $F(t, T_i)$ in (24.26), $F(t, T_i) + K$ is a geometric Brownian motion with adjusted volatility,

$$d(F(t, T_i) + K) = (F(t, T_i) + K) \frac{F(0, T_i)}{F(0, T_i) + K} \sigma(t, T_i) \cdot dW_{T_i}(t). \quad (24.29)$$

This allows us to apply Margrabe’s approach Margrabe (1978) for options to exchange one asset for another by rewriting the pay-off in (24.23) such that

$$\begin{aligned} C_{\text{Spread}}^{\text{Fwd}}(0, T_i, T_j, K) &= B(0, T_i) \mathbb{E}_{T_i} \left[(F(T_i, T_i) + K) \left(\frac{F(T_i, T_j)}{F(T_i, T_i) + K} - 1 \right)^+ \middle| \mathcal{F}_0 \right]. \end{aligned} \quad (24.30)$$

The solution of (24.29) is

$$\begin{aligned} &\frac{F(t, T_i) + K}{F(0, T_i) + K} \\ &= \exp \left\{ -\frac{1}{2} \left(\frac{F(0, T_i)}{F(0, T_i) + K} \right)^2 \int_0^t \sigma(s, T_i)^\top \sigma(s, T_i) ds \right. \\ &\quad \left. + \left(\frac{F(0, T_i)}{F(0, T_i) + K} \right) \int_0^t \sigma(s, T_i) \cdot dW_{T_i}(s) \right\}, \end{aligned}$$

where the right-hand side is an exponential martingale that defines a Radon–Nikodým derivative $d\hat{\mathbb{P}}/d\mathbb{P}$. Hence, the call price (24.30) can be written as

$$\begin{aligned} C_{\text{Spread}}^{\text{Fwd}}(0, T_i, T_j, K) \\ = B(0, T_i)(F(0, T_i) + K) \mathbb{E}_{\hat{\mathbb{P}}}[(Y(T_i) - 1)^+ | \mathcal{F}_0], \end{aligned} \quad (24.31)$$

where the dynamics of $Y(t) = F(t, T_j)/(F(t, T_i) + K)$,

$$\begin{aligned} dY(t) = Y(t) & \left[\sigma(t, T_j)^\top \Gamma_{i,j-1}(t) dt \right. \\ & \left. + \left(\sigma(t, T_j) - \frac{F(0, T_i)}{F(0, T_i) + K} \sigma(t, T_i) \right) \cdot dW_{\hat{\mathbb{P}}}(t) \right], \end{aligned}$$

is derived from (24.27) and (24.29) using Itô's product rule and

$$dW_{\hat{\mathbb{P}}}(t) = dW_{T_i}(t) - \frac{F(0, T_i)}{F(0, T_i) + K} \sigma(t, T_i) dt.$$

Standard techniques applied to (24.31) yield

$$\begin{aligned} C_{\text{Spread}}^{\text{Fwd}}(0, T_i, T_j, K) \\ = B(0, T_i) [F(0, T_j) e^{\int_0^{T_i} \sigma(s, T_j)^\top \Gamma_{i,j-1}(s) ds} N(d_+) \\ - (F(0, T_i) + K) N(d_-)], \end{aligned} \quad (24.32)$$

where $N(\cdot)$ is the cumulative standard normal distribution function and

$$\begin{aligned} d_{\pm} &= \frac{1}{\zeta} \left[\ln \left(\frac{F(0, T_j)}{F(0, T_i) + K} \right) + \int_0^{T_i} \sigma(s, T_j)^\top \Gamma_{i,j-1}(s) ds \pm \frac{1}{2} \zeta^2 \right], \\ \zeta^2 &= \int_0^{T_i} \left\| \sigma(s, T_j) - \frac{F(0, T_i)}{F(0, T_i) + K} \sigma(s, T_i) \right\|^2 ds. \end{aligned}$$

Figure 24.12 demonstrates the applicability of this approximation approach for the market scenario as calibrated in Section 24.6. The commodity forward maturities are chosen to be 1 year for T_i and 1 year plus 3 months for T_j . Since the forward interest rate maturities do not exactly match with the commodity forward maturities, we have interpolated the commodity forward prices in order to obtain values for $F(\cdot, T_i)$ and $F(\cdot, T_j)$. For an arbitrage-free interpolation in the maturities of the forward interest rates we refer to Schlögl (2002a). The resulting prices for call spread options on these forwards calculated by (24.32) and reference prices are shown for strikes between -15 and 10 in the left graph of Figure 24.12. Reference prices are computed by Monte Carlo simulation using Equations (24.24) and (24.25), where both commodity forwards as well as the interest forward rates $L(t, T_i)$ in (24.24) were simulated by the same set of 500,000 Brownian motion paths for all strikes.

The differences between the prices calculated by (24.32) and the Monte Carlo prices is shown by the upper curve of the right graph of Figure 24.12 and result from two approximations in the derivation of the formula. First, from ‘freezing’ the forward interest rates in (24.27) in order to make $\Gamma_{i,j-1}(t)$ deterministic and, secondly, from inserting the strike K in the dynamics of $F(t, T_i)$ as done in (24.29). The lower curve in the right graph shows the error resulting from the second approximation only. It is the difference between the prices calculated from (24.32) and Monte Carlo simulated prices from the dynamics given by (24.25) (i.e. without ‘freezing’ of the forward interest rates) and by (24.29).




FIGURE 24.12 Call spread options on forwards for the market as calibrated in Section 24.6. The maturities for the commodity forwards are 1 year and 1 year plus 3 months. Left: The black dots show Monte Carlo prices with simulated commodity forwards and simulated interest forward rates (i.e. no ‘freezing’). The blue solid line indicates the prices calculated using (24.32). Right: The black solid line shows the absolute difference between prices calculated by the closed-form formula in (24.32) and the Monte Carlo prices, i.e. the difference between the curves in the left graph. The dashed blue line shows the difference between the closed-form formula prices and a Monte Carlo simulated price using (24.27) and (24.29) but without ‘freezing’ in (24.28).

We conclude this section by sketching the derivation of the analogous spread call options formula for futures instead of forwards. From the forward dynamics (24.26) and relation (24.22) we obtain for futures the general dynamics

$$\begin{aligned} dG(t, T_k) = & -G(t, T_k)\sigma(t, T_k)^\top \bar{\eta}(t, T_k)dt \\ & + G(t, T_k)\sigma(t, T_k) \cdot dW_{T_k}(t) \quad (1 \leq k \leq N). \end{aligned} \quad (24.33)$$

Applying Kirk’s idea gives for the first futures with maturity T_i ,

$$\begin{aligned} d(G(t, T_i) + K) & = -(G(t, T_i) + K) \frac{G(0, T_i)}{G(0, T_i) + K} \sigma(t, T_i)^\top \bar{\eta}(t, T_i)dt \\ & + (G(t, T_i) + K) \frac{G(0, T_i)}{G(0, T_i) + K} \sigma(t, T_i) \cdot dW_{T_i}(t), \end{aligned} \quad (24.34)$$

which is the analogue to (24.29). For the second futures with maturity T_j we employ (24.5) in order to obtain the dynamics under dW_{T_i} ,

$$\begin{aligned} dG(t, T_j) = & G(t, T_j)\sigma(t, T_j)^\top [\Gamma_{i,j-1}(t) - \bar{\eta}(t, T_j)]dt \\ & + G(t, T_j)\sigma(t, T_j) \cdot dW_{T_i}(t). \end{aligned}$$

Using (24.21) we know that

$$\begin{aligned} \Gamma_{i,j-1}(t) & = \sum_{\ell=i}^{j-1} \frac{\delta L(0, T_\ell)}{1 + \delta L(0, \ell)} \lambda(t, T_\ell) \\ & = \sum_{\ell=i}^{j-1} \bar{\eta}(t, T_{\ell+1}) - \bar{\eta}(t, T_\ell), \end{aligned}$$

and the dynamics of the second futures is given by

$$\begin{aligned} dG(t, T_j) = & -G(t, T_j)\sigma(t, T_j)^\top \bar{\eta}(t, T_i)dt \\ & + G(t, T_j)\sigma(t, T_j) \cdot dW_{T_i}(t), \end{aligned}$$

which is the analogue to (24.27). Solving (24.34) yields the relation

$$\begin{aligned} & \frac{G(0, T_i) + K}{G(0, T_i) + K} \exp \left\{ \frac{G(0, T_i)}{G(0, T_i) + K} \int_0^t \sigma(s, T_i)^\top \bar{\eta}(s, T_i) ds \right\} \\ & = \exp \left\{ -\frac{1}{2} \left(\frac{G(0, T_i)}{G(0, T_i) + K} \right)^2 \int_0^t \sigma(s, T_i)^\top \sigma(s, T_i) ds \right. \\ & \quad \left. + \frac{G(0, T_i)}{G(0, T_i) + K} \int_0^t \sigma(s, T_i) \cdot dW_{T_i}(s) \right\}, \end{aligned}$$

in which the right-hand side is an exponential martingale and defines a Radon–Nikodým derivative $d\hat{\mathbb{P}}/d\mathbb{P}_{T_i}$. This allows us to write the spread call option prices as

$$\begin{aligned} & C_{\text{Spread}}^{\text{Fut}}(0, T_i, T_j, K) \\ & = B(0, T_i) \mathbb{E}_{T_i}[(G(T_i, T_j) - G(T_i, T_i) - K)^+ | \mathcal{F}_0] \\ & = B(0, T_i) \mathbb{E}_{T_i} \left[(G(T_i, T_i) + K) \left(\frac{G(T_i, T_j)}{G(T_i, T_i) + K} - 1 \right)^+ \middle| \mathcal{F}_0 \right] \\ & = B(0, T_i)(G(0, T_i) + K) \\ & \quad \times \exp \left\{ -\frac{G(0, T_i)}{G(0, T_i) + K} \int_0^{T_i} \sigma(s, T_i)^\top \bar{\eta}(s, T_i) ds \right\} \\ & \quad \times \mathbb{E}_{\hat{\mathbb{P}}}[(Y(T_i) - 1)^+ | \mathcal{F}_0], \end{aligned}$$

now with $Y(t) = G(t, T_j)/(G(t, T_i) + K)$. The same standard techniques as for spread options on forwards yield

$$\begin{aligned} & C_{\text{Spread}}^{\text{Fut}}(0, T_i, T_j, K) \\ & = B(0, T_i) \left[G(0, T_j) \exp \left\{ \int_0^{T_i} \sigma(s, T_j)^\top \bar{\eta}(s, T_i) ds \right\} N(d_+) \right. \\ & \quad \left. - (G(0, T_i) + K) \exp \left\{ -\frac{G(0, T_i)}{G(0, T_i) + K} \int_0^{T_i} \sigma(s, T_i)^\top \bar{\eta}(s, T_i) ds \right\} N(d_-) \right], \end{aligned}$$

with

$$\begin{aligned} d_\pm & = \frac{1}{\zeta} \left[\ln \left(\frac{G(0, T_j)}{G(0, T_i) + K} \right) \right. \\ & \quad \left. + \int_0^{T_i} (\sigma(s, T_j)^\top \bar{\eta}(s, T_i) \right. \\ & \quad \left. + \frac{G(0, T_i)}{G(0, T_i) + K} \sigma(s, T_i)^\top \bar{\eta}(s, T_i)) ds \pm \frac{1}{2} \zeta^2 \right], \\ \zeta^2 & = \int_0^{T_i} \left\| \sigma(s, T_j) - \frac{G(0, T_i)}{G(0, T_i) + K} \sigma(s, T_i) \right\|^2 ds. \end{aligned}$$

24.8 Conclusion

In the present paper, a joint model of commodity and interest rate dynamics based on the LIBOR Market Model approach was constructed. We have demonstrated how it can be effectively calibrated to market data, including at-the-money implied volatilities, and how less liquid instruments such as commodity spread options can be priced relative to the market using the model.

In closing, one should note that although neither seasonal cycles in commodity prices nor mean reversion of the commodity price process were explicitly considered in the model construction, both of these features (well documented in the empirical literature) can be captured by the model, to the extent that they are reflected in current market prices. Firstly, seasonal cycles in commodity prices are anticipated by the market and thus subsumed in the term structure of futures (or forward) prices, to which the model is calibrated by construction. Secondly, mean reversion is reflected in the term structure of volatility, to which the model is also calibrated, e.g. in the presence of mean reversion of the commodity spot price process, the volatility of the forward (or futures) prices increases as time to maturity decreases. Thus the martingale approach to the model construction does *not* preclude mean reversion simply because the objects considered explicitly are driftless—in fact, the corresponding processes under the appropriate probability measures are necessarily just as driftless in any other arbitrage-free model, including those that make mean reversion explicit in an Ornstein/Uhlenbeck process.

The main advantage of the LMM approach lies in its calibration to the market, and the model presented here opens up interesting avenues of further research (beyond the scope of the present paper) in terms of calibrating the model not just in the maturity dimension, but also in the strike dimension, along the lines of the ‘smile’-fitting extensions of the basic LMM discussed by Brace (2007).

Acknowledgements

This research was supported as part of a project with SIRCA under the Australian Research Council’s Linkage funding scheme (project No. LP0562616). The authors would like to thank the Editors of *Quantitative Finance* and two anonymous referees for helpful comments. The usual disclaimer applies.

References

- Black, F., The pricing of commodity contracts. *J. Financ. Econometr.*, 1976, **3**, 167–179.
- Brace, A., *Engineering BGM*, 2007 (Chapman & Hall/CRC: London).
- Brace, A., Gatarek, D. and Musiela, M., The market model of interest rate dynamics. *Math. Finance*, 1997, **7**, 127–154.
- Choy, B., Dun, T. and Schlögl, E., Correlating market models. *Risk*, 2004, **17**(9), 124–129.
- Cortazar, G. and Schwartz, E.S., The valuation of commodity contingent claims. *J. Deriv.*, 1994, **1**, 27–39.
- Cox, J.C., Ingersoll, J.E. and Ross, S.A., The relation between forward prices and futures prices. *J. Financ. Econometr.*, 1981, **9**, 321–346.
- Gibson, R. and Schwartz, E.S., Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance*, 1990, **45**, 959–976.
- Heath, D., Jarrow, R. and Morton, A., Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 1992, **60**, 77–105.
- Jamshidian, F., LIBOR and swap market models and measures. *Finance Stochast.*, 1997, **1**, 293–330.
- Kirk, E., Correlation in the energy markets. In *Managing Energy Price Risk*, edited by R. Jameson, 1995 (Risk Publications and Enron: London).

- Margrabe, W., The value of an option to exchange one asset for another. *J. Finance*, 1978, **33**, 177–186.
- Miltersen, K., Commodity price modelling that matches current observables: A new approach. *Quantit. Finance*, 2003, **3**, 51–58.
- Miltersen, K., Sandmann, K. and Sondermann, D., Closed form solutions for term structure derivatives with log-normal interest rates. *J. Finance*, 1997, **52**, 409–430.
- Miltersen, K. and Schwartz, E., Pricing of options on commodity futures with stochastic term structures of convenience yield and interest rates. *J. Financ. Quantit. Anal.*, 1998, **33**, 33–59.
- Musiela, M. and Rutkowski, M., Continuous-time term structure models: Forward measure approach. *Finance Stochast.*, 1997, **1**, 261–291.
- Pedersen, M.B., Calibrating Libor market models. SimCorp Financial Research Working Paper, 1998.
- Pelsser, A., *Efficient Methods for Valuing Interest Rate Derivatives*, 2000 (Springer: Berlin).
- Schlögl, E., Arbitrage-free interpolation in models of market observable interest rates. In *Advances in Finance and Stochastics*, edited by K. Sandmann and P. Schönbucher, 2002a (Springer: Heidelberg).
- Schlögl, E., A multicurrency extension of the lognormal interest rate market models. *Finance Stochast.*, 2002b, **6**, 173–196.
- Schwartz, E.S., The pricing of commodity-linked bonds. *J. Finance*, 1982, **37**, 525–539.
- Schwartz, E.S., The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance*, 1997, **52**, 923–973.

Appendix A: Volatility Integrals

This appendix describes the calculation of integrals as in (24.11) and (24.22) for piece-wise constant volatility matrices. We will focus on (24.22), since the integral in (24.11) can be obtained in the same way by setting $\bar{\eta} = \sigma$. The proposed method follows that used in Pedersen's LMM calibration (Pedersen 1998) to compute cap and swaption total variances.

In Section 24.3 a volatility matrix $U_C \in \mathbb{R}^{n_t \times d_C}$ was calculated for each calendar time $t_1, \dots, t_{n_t}$, and from the LMM calibration an analogous matrix $U_I \in \mathbb{R}^{m_t \times d_I}$ is available for each calendar time $s_1, \dots, s_{m_t}$. Merging over calendar times yields in each case three-dimensional arrays $\Lambda_C = (\lambda_{i,j,k}^C) \in \mathbb{R}^{n_t \times n_t \times d}$ and $\Lambda_I = (\lambda_{i,j,k}^I) \in \mathbb{R}^{m_t \times m_t \times d_I}$, respectively. Forward times are given by $x_1, \dots, x_{n_t}$ for commodity forwards and by $y_1, \dots, y_{m_t}$ for interest forward rates. The calendar times $t_0 = s_0 = 0$ and forward times $x_0 = y_0 = 0$ are added for notational convenience. From these matrices of piecewise constant volatilities the integral in Equation (24.22) has to be calculated.

The first step illustrates the computation of $\bar{\eta}$ and can be omitted when calculating the integral in (24.11). The matrix Λ_I determines the forward interest rate volatilities $\lambda(s, s + y)$ as occurring in Equation (24.6). From these, a corresponding matrix $\Theta_I \in \mathbb{R}^{m_t \times m_t \times d_I}$ with entries $\bar{\eta}_{i,j,k}$ is calculated using relation (24.18). Unfortunately, the forward rates $L(t, T)$ in (24.18) are not known for $t > s_0$, but as a best guess we employ the actual forward rate curve in the sense that, for all future times $s_i > s_0$, an approximated forward rate curve $\bar{L}(s_i, T)$ is given by $L(s_0, T)$, essentially again making use of the standard 'frozen coefficient' approximation.

In the second step, the integral for a given time to maturity T in the upper integration limit is computed. We equalize the calendar time discretization of the volatility matrices for commodity forwards and forward interest rates by taking the union of $t_0, \dots, t_{n_t}$ and $s_0, \dots, s_{m_t}$. The new calendar times will be denoted by $t_0 = 0, t_1, \dots, t_{p_t}$. The same is done for the discretization of times to maturity, which results in a discretization $x_0 = 0, x_1, \dots, x_{p_t}$. The arrays Λ_C and Θ_I are extended accordingly (i.e. by simply inserting lines for different calendar times and columns for different times to maturity, respectively, but otherwise identical entries $\lambda_{i,j,k}$ and $\bar{\eta}_{i,j,k}$), such that $\Lambda_C, \Theta_I \in \mathbb{R}^{p_t \times p_t \times d}$. The purpose of this extension is to have a twofold sum in formula (A.1) instead of a fourfold sum.

Any tuple in $\mathcal{I} = \{(i, j) : 1 \leq i \leq p_c, 1 \leq j \leq p_f\}$ relates to a pair of $\sigma(t_i, t_i + x_j)$ and $\bar{\eta}(t_i, t_i + x_j)$, and in order to be relevant for the considered integral, the following inequalities need to be satisfied for some $0 \leq t \leq T$:

$$t_{i-1} \leq t < t_i, \quad x_{j-1} \leq T - t < x_j.$$

Equivalently, this can be written as

$$\kappa_{i,j} := \min\{t_i, T - x_{j-1}, T\} - \max\{t_{i-1}, T - x_j\} > 0.$$

Hence, the integral is given by

$$\int_t^T \sigma(u, T) \bar{\eta}(u, T) du = \sum_{(i,j) \in \mathcal{J}} \kappa_{i,j} \sum_{k=1}^d \lambda_{i,j,k}^C \bar{\eta}_{i,j,k}, \quad (\text{A.1})$$

where $\mathcal{J} = \{(i, j) \in \mathcal{I} : \kappa_{i,j} > 0\}$.

IV

Electricity Markets

25	Modelling the Distribution of Day-Ahead Electricity Returns: A Comparison <i>Alessandro Sapio</i>	551
	Introduction • Data and Preliminary Analysis • Distributions of Electricity Returns • Fitting the Empirical Probability Densities • Robustness • Conclusion	
26	Stochastic Spot Price Multi-Period Model and Option Valuation for Electrical Markets <i>Eivind Helland, Timur Aka, and Eric Winnington</i>	571
	Introduction • Spot Price Multi-Period Model • Model Calibration • Scenario Generation and Valuation of Structured Products • Structured Product Valuation • Conclusion	
27	Modelling Spikes and Pricing Swing Options in Electricity Markets <i>Ben Hambly, Sam Howison, and Tino Kluge</i>	585
	Introduction • Properties of the Model for Spot Prices • Option Pricing • Pricing Swing Options	
28	Efficient Pricing of Swing Options in Lévy-Driven Models <i>Oleg Kudryavtsev and Antonino Zanette</i>	607
	Introduction • Lévy Processes: Basic Facts • The Multiple Optimal Stopping Problem for Swing Options • The Finite Difference Scheme for Pricing Swing Options • Pricing Swing Options Using the Wiener–Hopf Approach • Numerical Results	
29	Hedging Strategies for Energy Derivatives <i>Peter Leoni, Nele Vandaele, and Michèle Vanmaele</i>	623
	Introduction • Literature • LRM Hedging Strategy • Application to Non-Traded Assets • Numerical Results • Conclusion	
30	The Valuation of Clean Spread Options: Linking Electricity, Emissions and Fuels <i>René Carmona, Michael Coulon, and Daniel Schwarz</i>	645
	Introduction • The Bid Stack: Price Setting in Electricity Markets • Risk-Neutral Pricing of Allowance Certificates • Valuing Clean Spread Options • A Concrete Two-Fuel Model • Numerical Analysis	
31	Is the EUA a New Asset Class? <i>Vicente Medina and Angel Pardo</i>	667
	Introduction • Stylised Facts of Asset Returns • Data • Results • Conclusions	

Modelling the Distribution of Day-Ahead Electricity Returns: A Comparison

25.1	Introduction	551
25.2	Data and Preliminary Analysis	553
25.3	Distributions of Electricity Returns.....	556
25.4	Fitting the Empirical Probability Densities	558
25.5	Robustness	562
25.6	Conclusion	565
	References.....	568

Alessandro Sapio

This paper contributes to the characterization of the probability density of the price returns in some European day-ahead electricity markets (NordPool, APX, Powernext) by fitting flexible and general families of distributions, such as the α -stable, Normal Inverse Gaussian (NIG), Exponential Power (EP) and Asymmetric Exponential Power (AEP) distributions, and comparing their goodness of fit. The α -stable and the NIG systematically outperform the EP and AEP models, but the tail behaviour and the skewness estimates are sensitive to the definition of the returns and to the deseasonalization methods. In particular, the logarithmic transform and volatility rescaling tend to dampen the extreme returns.

Keywords: Energy markets; Non-Gaussian distributions; Energy economics; Empirical time series analysis

25.1 Introduction

The ‘problem of price variation’, as Mandelbrot (1963) called it, has been one of the most debated issues in financial economics. Providing a correct description of the empirical distribution of returns is essential for the theory and practice of trading and investment. Indeed, portfolio selection theory based upon variance-based measures of risk only works under the assumption that returns have finite central moments. Furthermore, Value-at-Risk calculations, the pricing formulas for contingent claims, the accuracy of price forecasting and the appropriateness of econometric methods all depend on the distribution of returns.

The relevance of these issues is by no means confined to stock market analysis. In markets for non-storable commodities, such as electricity, trading mechanisms must continuously ensure market-clearing. In electricity, imbalances between demand and supply would cause blackouts. The prices quoted on wholesale power exchanges undergo sudden and short-lived excursions, caused by

strategic behaviour and accidental plant failures, while production and consumption smoothing are not feasible. A well-known empirical fact is that the heavy tails observed in electricity return distributions cannot be accounted for by the Gaussian law. Forecasting and risk management are therefore even more crucial than in stock markets, and a 'solution' to the problem of price variation in the context of power exchanges is even more urgent.

It is the aim of this paper to investigate the distributional nature of the day-ahead electricity price returns. More specifically, one asks whether electricity returns display heavy tails and skewness, and whether the central moments diverge; one also investigates the time scaling of risk and possible intraday differences in distributional shapes. Answers to these research questions are sought by comparing the goodness-of-fit performances of the α -stable, the Normal Inverse Gaussian (NIG), the Exponential Power (EP) and the Asymmetric Exponential Power (AEP) distribution laws. The data are drawn from major European power exchanges, such as the Scandinavian NordPool, the Dutch APX and the French Powernext. We focus on one-day returns computed on prices of individual hours, which allow us to obtain a grasp on the intraday risk patterns.

The main findings of the study are the following. First, and expectedly, electricity returns display heavy tails. This fact is robust across markets and hourly auctions, and holds for various definitions of returns (log-returns, percentage returns, price changes) and regardless of the deseasonalization methodology. The tails in the APX market and in day-time auctions are fatter than in the other markets and in night-time auctions, and tend to be dampened by the logarithmic transform. Second, the NIG and the α -stable laws systematically outperform the EP and AEP distributions according to goodness-of-fit criteria, although no clear ranking can be established between the two 'winners'. Because only the α -stable and the NIG distributions are closed under convolution, fat tails are also expected to characterize returns computed over longer horizons. The estimated characteristic exponent of the α -stable distribution is always between 1 and 2, indicating that the expected value of the electricity returns converges, but the second moment does not. However, this is true only to the extent that the α -stable outperforms the NIG. Third, the skewness is an essential feature of the returns distributions. Some interesting cross-market variance emerges when Cholesky/scaling log-returns are considered (negative skewness in Powernext, positive in APX). Yet, using other definitions of returns and deseasonalization methods yields mixed results.

A number of previous works are closely related to the present paper. The α -stable model was shown to outperform the Hyperbolic and NIG distributions by Rachev *et al.* (2004, cited by Weron 2009) for EEX daily price differences, and by Weron (2005) for EEX and NordPool data. Mugele *et al.* (2005) found that NordPool and EEX daily price differences were also best described by stable laws, while the performance of the stable distribution for PolPX data was less successful. Further evidence of power-law tails was found by Bellini (2002), Byström (2005) and Chan and Gray (2006), who fitted generalized extreme-value distributions by means of peaks-over-threshold and block maxima methods, and by Deng and Jiang (2005), who followed a quantile function approach to model the distribution of CalPX and PJM returns. The Generalized Hyperbolic distribution, which includes the NIG as a special case, was estimated on NordPool data by Eberlein and Stahl (2003). Further work by Weron (2009) reports estimates of the α -stable, Hyperbolic and NIG models on data from several markets (EEX, Omel, PJM, NEPOOL) and using various measures of price returns. The results vary across countries and are affected by how the returns are defined (see also Weron 2006). The probability density function of daily log-returns was modelled as a symmetric EP distribution by Bottazzi *et al.* (2005) and Bottazzi and Sapio (2007), whose estimates hinted at Laplacian or even heavier tails. Robinson and Baniak (2002) also fitted a Laplace distribution, whereas Bosco *et al.* (2007) used EP-distributed shocks in a PARMA-GARCH model. Deng *et al.* (2002) fitted a Cauchy-Laplace mixture to PJM and CalPX price differences. The present study covers three markets, analyses relatively large samples, and considers three definitions of price returns (log-returns, percentage returns, price differences). In this respect, the paper seeks to overcome the main limitations of the previous works, as highlighted, for instance, by Weron (2009, p. 460).

The paper is organized as follows. Section 25.2 describes the datasets and the deseasonalization methods, and provides summary statistics. The distributional models fitted in the paper are described in Section 25.3, whereas Section 25.4 illustrates the baseline estimation results concerning log-returns. In Section 25.5, the robustness of the results is assessed with respect to other definitions of returns and deseasonalization methods. Section 25.6 concludes.

25.2 Data and Preliminary Analysis

For the purposes of this study, data on day-ahead electricity prices were collected for three major European power exchanges: NordPool (Denmark, Finland, Norway, Sweden), 2191 days from 1 January 1997 to 31 December 2002; APX (the Netherlands), 1457 days from 6 January 2001 to 31 December 2004; and Powernext (France), 1826 days from 1 February 2002 to 31 January 2007.* In these markets, each day, 24 auctions are run simultaneously in order to determine prices and quantities for each hour of the following day. The day-ahead prices are determined by means of uniform price auctions, so that all power is sold and purchased at the market-clearing price. The time series of day-ahead prices are depicted in Figure 25.1 for the 4 p.m. auctions, when demand is typically near its daily peak, and in Figure 25.2 for a night delivery session (4 a.m.), when average prices and demand are relatively low.

In finance, the price returns are usually defined as logarithmic price differences or log-returns $x_{ht} = \log p_{ht} - \log p_{h,t-1}$, where p_{ht} is the price at day t for the hour- h auction.† Table 25.1 provides summary statistics of the log-returns for selected hours, along with the outcomes of Shapiro–Wilk normality tests




FIGURE 25.1 Plots of NordPool, APX and Powernext day-ahead prices for the 4 p.m. auctions. NOK, Norwegian kroner; Eur, Euro; kWh, KiloWatts per hour.

* Data sources: NordPool FTP server, www.apx.nl and www.powernext.fr, respectively.

† Later in the paper (Section 25.5) we shall discuss some drawbacks when using log-returns and assess the robustness of the results with respect to other definitions of returns.




FIGURE 25.2 Plots of NordPool, APX and Powernext day-ahead prices for the 4 a.m. auctions. NOK, Norwegian kroner; Eur, Euro; KWh, KiloWatts per hour.

TABLE 25.1 Summary Statistics of Log-Returns in the NordPool, APX and Powernext Markets, Along with Shapiro–Wilk Statistics and Autocorrelation Coefficients

Auction	Mean	Std. Dev.	Skewness	Kurtosis	SW	acf(1)	acf(7)
<i>NordPool</i>							
4 a.m.	−0.0000	0.1380	0.4668	28.2410	0.7167	−0.2715	0.0371
8 a.m.	−0.0000	0.2107	1.0293	19.9273	0.7829	−0.1508	0.5278
12 (noon)	−0.0000	0.1384	1.0630	20.5106	0.8204	−0.1233	0.4812
4 p.m.	−0.0000	0.1193	1.2205	14.4644	0.8291	−0.0384	0.5035
8 p.m.	0.0000	0.0936	0.8641	43.0378	0.7597	−0.1489	0.1823
12 (midnight)	0.0000	0.0623	−0.3141	13.1957	0.8584	−0.0684	0.1075
<i>APX</i>							
4 a.m.	−0.0009	1.2534	0.2439	23.7528	0.5273	−0.4427	0.0204
8 a.m.	−0.0002	1.6831	0.3925	15.1495	0.6517	−0.3893	0.3665
12 (noon)	0.0003	0.6921	0.8037	6.5865	0.9355	−0.2268	0.5297
4 p.m.	−0.0004	0.6345	0.8864	9.8323	0.8842	−0.1865	0.4836
8 p.m.	−0.0003	0.3410	−0.1048	9.8369	0.8879	−0.3801	0.1472
12 (midnight)	−0.0002	0.4831	0.6416	169.4452	0.4380	−0.4664	0.0193
<i>Powernext</i>							
4 a.m.	−0.0000	0.5119	0.1385	83.7746	0.6818	−0.3742	0.3586
8 a.m.	−0.0000	0.6526	0.7423	6.8751	0.9071	−0.2713	0.6393
12 (noon)	−0.0000	0.4424	1.0837	11.1308	0.8926	−0.2023	0.4539
4 p.m.	0.0000	0.4484	1.0351	9.3050	0.9049	−0.2118	0.5202
8 p.m.	−0.0000	0.2924	0.6311	6.7971	0.9283	−0.1523	0.3808
12 (midnight)	0.0000	0.2103	0.4558	13.4018	0.8779	−0.2977	0.1724

and autocorrelation coefficients. This table shows that, while drifts in power prices are rather weak, the standard deviations are largest at the beginning of the working day (8 a.m.) and decay thereafter. The skewness is positive and stronger during the day, and the excess kurtosis is always positive and large, albeit without any clear intraday pattern. Hence, the probability of observing large positive or negative fluctuations is greater than in a Gaussian process. The Shapiro–Wilk normality tests strongly reject the null of a Gaussian distribution ($SW = 1$) for all markets and all hourly auctions: the test statistics are always significantly below 1 (p values, not reported here, are all below 0.0001).

The serial correlations over a daily horizon are always negative, more so in the night-time auctions; the autocorrelations at lag 7 days are strong, up to 0.5–0.6 in some day-time auctions. The weekly pattern of economic activity is an obvious determinant of these patterns. Further time dependencies appear at lower frequencies, due to the seasonal patterns of economic activity and weather conditions. In addition to linear dependencies, the width of the power price fluctuations vary across hours, because bidding strategies change under different relative scarcity (Karakatsani and Bunn 2004, Bottazzi *et al.* 2005, Simonsen 2005). All of this justifies the use of filters in order to remove the linear and higher-order autocorrelations in such a way that what remains is presumably the outcome of random shocks to market fundamentals.

The data are filtered in two steps. First, we remove all the linear autocorrelations by means of the semiparametric Cholesky factor algorithm introduced by Diebold *et al.* (1997). The algorithm works as follows.

1. Estimate the covariance matrix Σ of the vector x_{ht} as the Toeplitz matrix built upon the autocovariance vector γ^* .
2. Calculate C as the Cholesky factor of Σ , i.e. $C: CC' = \Sigma$.
3. Extract the linearly uncorrelated, standardized residuals $\tilde{x}_{ht}$ as follows:

$$\tilde{x}_{ht} = C^{-1}x_{ht}. \quad (25.1)$$

Let us call $\tilde{x}$ the *Cholesky-filtered* log-returns. As a second step, we model the standard deviation of the filtered returns as a power function of the lagged price level (which is a proxy for market scarcity) or, in logs,

$$\log V[\tilde{x}_{ht} | p_{h,t-1}] = \chi + \chi' d_{ht} + \rho \log p_{h,t-1} + \rho' (\log p_{h,t-1}) d_{ht} + \varepsilon_{ht}, \quad (25.2)$$

and rescale the filtered returns in order to obtain homoskedastic samples[†]:

$$x_{ht}^* = \frac{x_{ht}}{e^{\hat{\chi} + \hat{\chi}' d_{ht} + \hat{\rho} \log p_{h,t-1} + \hat{\rho}' (\log p_{h,t-1}) d_{ht}}}. \quad (25.3)$$

Finally, the hourly averages are subtracted. These returns will be referred to as *Cholesky-filtered and rescaled* log-returns. In the above equations, $V[\cdot]$ is the variance operator, χ , χ' , ρ and ρ' are constant coefficients (their estimated values are indicated with a hat), $\tilde{x}_{ht}$ is the Cholesky-filtered series of log-returns for the hour- h auction at day t , $p_{h,t-1}$ is the price for the hour- h auction at day $t-1$, and ε_{ht} is an i.i.d. error term. The dummy variable d_{ht} allows both the slope and the intercept of the scaling regression to vary as the price reaches particularly high levels. This accounts for the possibility that the price dynamics is characterized by switching regimes (see De Jong (2006) and Weron (2009) and references therein). In order to estimate the power-law scaling coefficients, the data of each time series are grouped into equipopulation bins. Next, the sample standard deviations of the log-returns in each bin are computed, and

* A Toeplitz matrix is a matrix that has constant values along all negative-sloping diagonals.

† A scaling relationship between log-return variance and volume levels was also considered, but it is seldom significant.

the logarithm of the sample standard deviations is OLS-regressed on a constant and on the logarithm of the mean price level within the corresponding bins.*

The estimates of the variance–price scaling relationship for NordPool and Powernext suggest that the standard deviations of the filtered returns are negatively correlated with the lagged price levels, but the dummy coefficients are large and positive, implying that the variance–price relationship is increasing at high price levels. The APX scaling exponents, instead, are quite variable across hours. Ljung–Box tests performed using 28 lags (4 weeks) cannot reject the null of zero serial correlation, while normality tests still reject the null of a Gaussian distribution.† Even after filtering and rescaling, the log-returns still display skewness and excess kurtosis.

This filtering procedure departs from what Weron (2009) calls the ‘industry standard’, wherein the electricity price is envisaged as the sum or the product of a (deterministic) trend/cycle/seasonal component and a stochastic component. The goal of the deseasonalization techniques used by De Jong (2006) and Weron (2009) is to isolate the stochastic component, which is then used to compute the price returns. De Jong (2006) regresses the log-prices on daily dummies, an annual sinusoidal, and an exponentially weighted moving average, while Weron (2009) applies a wavelet smoothing technique to deal with the annual cycle, plus a moving-average filter to remove the average weekly pattern. A problem with these approaches is that they may not yield the desirable i.i.d. samples. The approach followed in this paper is a way to solve this problem: the outcome of the Cholesky filter is a serially uncorrelated time series by construction, while the variance–price power-law scaling takes care of the remaining heteroskedasticity.

25.3 Distributions of Electricity Returns

When characterizing the probability density of the electricity price returns, it is desirable to select classes of probability distributions that are general and flexible enough so as to yield different implications concerning the decay of the tails, the skewness, the convergence of the central moments and the time scaling of risk. In this work, we focus on the α -stable, Normal Inverse Gaussian, Exponential Power and Asymmetric Exponential Power distribution families. These have frequently been analysed in the relevant literature, as mentioned in the Introduction.

The first class of probability distributions is relevant if one views the electricity returns as resulting from the sum of n i.i.d. shocks u_j , not restricted to having finite moments, with $j = 1, \dots, n$. The Generalized Central Limit Theorem states that the distribution of $(1/\sqrt{n})\sum_{j=1}^n u_{jt}$ converges to an α -stable distribution as $n \rightarrow \infty$ (Samorodnitsky and Taqqu 1994, Borak *et al.* 2005). A random variable is α -stable if and only if its characteristic function reads

$$\phi(\nu) = e^{-\sigma^\alpha |\nu|^\alpha \{1 + i\beta(\text{sign}\nu) \tan(\pi\nu/2)[(|\sigma|\nu|)^{1-\alpha} - 1]\} + i\mu\nu} \quad (25.4)$$

if $\alpha \neq 1$, or

$$\phi(\nu) = e^{-\sigma|\nu|\{1 + i\beta(\text{sign}\nu)(2/\pi)\log(\sigma|\nu|)\} + i\mu\nu} \quad (25.5)$$

* Estimation of the scaling coefficients was performed for numbers of bins between 8 and 40. One finds that R^2 values are decreasing in the number of bins, and that the point estimates of $\hat{\rho}$ tend to decrease slightly in absolute value. A decision was made to focus on scaling based on 40 bins (NordPool), 28 bins (APX) and 34 bins (Powernext), corresponding to between 52 and 55 observations per bin. Indeed, Monte Carlo simulations performed by the author show that the profile of the scaling exponent estimates, with respect to the number of bins, is characterized by a flat region around the mentioned values. A larger number of bins implies more degrees of freedom in the regression, but the variance estimates within each bin are more noisy, because they are based on a smaller number of observations, resulting in smaller R^2 values. With respect to the use of the mean prices, choosing the median prices does not affect the results significantly.

† The summary statistics for the Cholesky-filtered and rescaled variables and detailed information on the variance–price scaling estimates are available upon request.

if $\alpha = 1$. An α -stable distribution is defined by four parameters: a characteristic exponent or stability index $\alpha \in (0, 2]$, a skewness parameter $\beta \in [-1, 1]$, a scale parameter $\sigma > 0$, and a location parameter $\mu \in \mathfrak{R}$. The stable distribution corresponds to a Normal when $\alpha = 2$, whereas $\alpha < 2$ implies that the variance is infinite and the tails asymptotically decay as power laws. When $\alpha = 1$, the Cauchy distribution results, but if $\alpha < 1$ even the first central moment diverges.

A second model assumes that the electricity returns are variance–mean mixtures of Gaussian random variables. If the mixing distribution is a generalized inverse Gaussian law with $\lambda = -1/2$, the Normal Inverse Gaussian (NIG) law obtains. The probability density function reads (Barndorff-Nielsen 1997)

$$f_{\text{NIG}}(x; \alpha, \beta, \sigma, \mu) = \frac{\alpha \sigma}{\pi} e^{\sigma \sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)} \frac{K_1(\alpha \sqrt{\sigma^2 + (x - \mu)^2})}{\sqrt{\sigma^2 + (x - \mu)^2}}. \quad (25.6)$$

The parameters are α (steepness), β (skewness), $\sigma > 0$ (scale) and $\mu \in \mathfrak{R}$ (location). The constant K_1 is the modified Bessel function of the third kind with index 1, also known as the MacDonald function. The NIG exhibits semi-heavy tails, i.e. heavier than Gaussian, but lighter than power law. The Cauchy distribution is the special case $f_{\text{NIG}}(x; 0, 0, 1, 0)$. The tails of the NIG distribution taper off according to the following asymptotic formula:

$$f_{\text{NIG}}(x) \approx |x|^{-3/2} e^{(\mp \alpha + \beta)x}, \quad (25.7)$$

for $x \rightarrow \pm\infty$. Note that both the α -stable and the NIG distribution are closed to convolution. This feature is particularly useful in the time scaling of risk, e.g. in deriving long-term risk from daily risk (Weron 2004).

The third class of distributions obtains if one assumes $x \sim \text{i.i.d. } N(0, h^\psi)$, where $h \sim \text{i.i.d. Exponential}$. As shown by Fu *et al.* (2005), if $\psi \geq 0$, this model yields an Exponential Power distribution with shape parameter $b > 0$ (inversely related to ψ), scale parameter a and position parameter μ :

$$f_{\text{EP}}(x; a, b, \mu) = \frac{1}{2ab^{1/b}\Gamma[1 + (1/b)]} e^{-(1/b)|(x - \mu)/a|^b}, \quad (25.8)$$

where $\Gamma(\cdot)$ is the gamma function.* The EP distribution reduces to a Laplace if $b = 1$ and to a Normal if $b = 2$. As b becomes smaller, the density becomes heavier-tailed and more sharply peaked.†

The EP family only includes symmetric probability distributions. However, it may be desirable to allow for more flexibility in modelling the skewness, in order to yield a more punctual comparison with the α -stable and NIG laws. The EP family has been generalized in this direction by Bottazzi and Secchi (2007), who introduced the Asymmetric Exponential Power (AEP) family:

$$\begin{aligned} f_{\text{AEP}}(x; a_l, a_r, b_l, b_r, \mu) \\ = \frac{1}{a_l A_0(b_l) + a_r A_0(b_r)} \\ \times e^{-(1/b_l)|(x - \mu)/a_l|^{b_l} \theta(\mu - x) + (1/b_r)|(x - \mu)/a_r|^{b_r} \theta(x - \mu)} \end{aligned} \quad (25.9)$$

where $\theta(y)$ (for a generic variable y) is the Heaviside theta function, and

$$A_k(y) = y^{[(k+1)/y]-1} \Gamma\left(\frac{k+1}{y}\right).$$

* This distribution was first used in economics by Bottazzi and Secchi (2003), and is also known as the Subbotin distribution (Subbotin 1923).

† West (1987) represented the Exponential Power distribution as a scale mixture of Normals with an α -stable mixing distribution whose stability index is equal to $b/2$, but his result is limited to $b \geq 1$. See also Andrews and Mellows (1974) and Choy and Walker (2003).

The AEP density is characterized by two positive shape parameters (b_l, b_r), two positive scale parameters (a_l, a_r), and one position parameter (μ). The magnitudes of the shape parameters tune the behaviour of the upper and lower tail, respectively. The AEP reduces to the EP distribution when $a_l = a_r$ and $b_l = b_r$. Unlike the α -stable and NIG distributions, the EP and AEP distributions are not closed under convolution. Hence, the distribution of returns computed over longer time horizons tends to converge to the Gaussian law.

The parameters of the stable distribution are estimated here by means of the characteristic function regression method (Koutrouvelis 1980, Kogon and Williams 1998). This method was shown by Weron (2004) and Scalas and Kim (2007) to be more accurate than alternative methods. We use Maximum Likelihood to estimate the parameters of the NIG, EP and AEP distributions. The MFE Toolbox (Weron 2006) is exploited for the estimation of stable and NIG laws, whereas the EP and AEP distributions were fitted by making use of the Subbotools package (see also Bottazzi 2004 and Bottazzi and Secchi 2007).*

The goodness of fit of the estimated distribution models is assessed by means of the Kolmogorov–Smirnov and Cramer–von Mises statistics (D’Agostino and Stephens 1986). The Kolmogorov–Smirnov D statistic is defined as the maximum absolute deviation between the theoretical and empirical CDFs:

$$D = \max \left(\left| \frac{i}{N} - z_i \right| \right), \quad (25.10)$$

where z_i is the i th ordinate of the theoretical cumulative distribution function under test, and N the sample size. The Cramer–von Mises W^2 test statistics is based on the quadratic deviations between theoretical and empirical CDFs:

$$W^2 = \frac{1}{12N} + \sum_{i=1}^N \left[z_i - \frac{2i-1}{2N} \right]^2. \quad (25.11)$$

The asymptotic 5% limiting value of D is $1.36 / \sqrt{N}$, that is 0.0291 (NordPool), 0.0356 (APX) and 0.0318 (Powernext). These are the Monte Carlo asymptotic values, under the assumption that sample sizes are large enough as to rule out the need for distribution-specific small sample corrections. The 5% limiting value for W^2 is 0.443—an exact result valid for any sample size greater than or equal to 5 (Stephens 1974).

25.4 Fitting the Empirical Probability Densities

The estimation results for the α -stable, NIG, EP and AEP distributions are reported in Tables 25.2–25.4 for selected hourly auctions in NordPool, APX and Powernext, respectively, along with goodness-of-fit statistics.

The estimated stability index α for the stable distribution is always below the Normal value (i.e. 2). NordPool and Powernext point estimates of α lie within the range 1.70–1.85; APX estimates are slightly lower, more so in the 4 a.m. and 8 a.m. auctions: therefore, the tails of APX log-returns decay more slowly. The skewness parameter β in APX is positive for all hours, whereas $\beta < 0$ for all NordPool and Powernext hourly auctions. These patterns are confirmed, at least qualitatively, by the estimated NIG parameters: the log-returns in all markets display heavy tails, the steepness parameter is lower in the APX market, and the distribution of Powernext log-returns is characterized by a negative skew in all hours. This time the parameter β in the NordPool assumes negative values only for some hours (4 a.m., midnight), and is below zero in some APX auctions (4 a.m., 8 a.m.).

As to the EP distribution, the point estimates of the shape parameter are systematically below the Normal value, a sign of heavy tails. The shape coefficients are often around the Laplace value (namely 1), with some deviations (most frequently above 1 in the Powernext and below 1 in the APX auctions). The AEP results show that the left tails are heavier in the NordPool and APX night auctions, and lighter during the day ($b_l < b_r$ by night, the opposite in the other hours). Consistent with the α -stable and NIG

* Subbotools is available at <http://cafim.sssup.it/~giulio/software/subbotools/>.

TABLE 25.2 Parameter Estimates and Goodness-of-fit Statistics for α -Stable, NIG, EP and AEP Distributions Fitted to the Filtered and Rescaled Log>Returns: NordPool

Auction	Distribution	Tail	Parameter		Location	Test value	
			Skewness	Scale		D	$W2$
4 a.m.	α -Stable	1.7750	-0.1536	0.5910	0.0263	0.0133	0.1908
	NIG	1.0299	-0.1707	1.0252	0.1723	0.0185	0.1461
	EP	1.1238		0.7803	0.0000	0.0380	0.6938
	AEP	1.0261, 1.4123		0.8033, 0.7961	0.0000	0.0300	0.3132
8 a.m.	α -Stable	1.6997	-0.0989	0.6101	-0.0177	0.0135	0.1066
	NIG	0.7107	0.0030	0.9164	-0.0039	0.0144	0.0655
	EP	0.9979		0.7992	0.0000	0.0211	0.2575
	AEP	1.0434, 0.9763		0.8014, 0.8047	0.0000	0.0249	0.2751
12 (noon)	α -Stable	1.7494	-0.0027	0.6447	-0.0162	0.0112	0.0991
	NIG	0.8139	0.0410	1.0472	-0.0528	0.0118	0.0479
	EP	1.0845		0.8471	0.0000	0.0263	0.2472
	AEP	1.1774, 1.0393		0.8512, 0.8562	0.0000	0.0220	0.2098
4 p.m.	α -Stable	1.7553	-0.1158	0.6350	-0.0213	0.0133	0.1711
	NIG	0.8327	0.0284	1.0324	-0.0352	0.0116	0.0470
	EP	1.0872		0.8324	0.0000	0.0256	0.2483
	AEP	1.1647, 1.0458		0.8350, 0.8391	0.0000	0.0250	0.2214
8 p.m.	α -Stable	1.7608	-0.1242	0.6446	-0.0285	0.0110	0.1276
	NIG	0.8325	0.0168	1.0559	-0.0213	0.0141	0.0583
	EP	1.0831		0.8411	0.0000	0.0235	0.2328
	AEP	1.1424, 1.0513		0.8432, 0.8467	0.0000	0.0234	0.2424
12 (midnight)	α -Stable	1.7946	-0.3804	0.6329	-0.0135	0.0202	0.2138
	NIG	1.0763	-0.1499	1.1719	0.1648	0.0120	0.0264
	EP	1.2292		0.8454	0.0000	0.0266	0.3865
	AEP	1.1209, 1.3770		0.8501, 0.8458	0.0000	0.0118	0.0487

Tail: α (stable and NIG), b (EP), b_l and b_r (AEP). Skewness: σ (stable and NIG). Scale: α (stable and NIG), a (EP), a_l and a_r (AEP). Location: μ . Asymptotic 5% limiting values: 0.0291 (D), 0.443 (W^2). Data in bold type indicate the minimum goodness-of-fit test values.

TABLE 25.3 Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions Fitted to the Filtered and Rescaled Log>Returns: APX

Auction	Distribution	Tail	Parameter		Location	Test value	
			Skewness	Scale		D	$W2$
4 a.m.	α -Stable	1.5158	0.3067	0.5181	0.2085	0.0219	0.0355
	NIG	0.4264	-0.0869	0.6303	0.1312	0.0348	0.3584
	EP	0.7562		0.6841	0.0000	0.0659	1.9361
	AEP	0.6904, 1.0266		0.7324, 0.7045	0.0000	0.0377	0.4536
8 a.m.	α -Stable	1.3876	0.3903	0.5054	0.3332	0.0341	0.0750
	NIG	0.3093	-0.0649	0.5851	0.1255	0.0355	0.3809
	EP	0.6893		0.6893	0.0000	0.0735	2.4395
	AEP	0.6231, 0.9073		0.7363, 0.7028	0.0000	0.0406	0.3575
12 (noon)	α -Stable	1.7749	0.6614	0.6317	0.0215	0.0215	0.2227
	NIG	1.1050	0.2664	1.1675	-0.2900	0.0132	0.0434

(Continued)

TABLE 25.3 (Continued) Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions Fitted to the Filtered and Rescaled Log-Returns: APX

Auction	Distribution	Tail	Parameter		Location	Test Value	
			Skewness	Scale		D	W^2
4 p.m.	EP	1.2244		0.8559	0.0000	0.0480	0.6441
	AEP	1.5558, 1.0607		0.8654, 0.8717	0.0000	0.0149	0.0565
	α -Stable	1.6632	0.1908	0.5824	-0.0119	0.0202	0.1229
	NIG	0.7513	0.1063	0.8877	-0.1269	0.0186	0.0812
	EP	1.0219		0.7840	0.0000	0.0429	0.4586
8 p.m.	AEP	1.1819, 0.9421		0.7880, 0.7993	0.0000	0.0213	0.1735
	α -Stable	1.6520	0.1309	0.5986	-0.0158	0.0161	0.4738
	NIG	0.7178	0.0784	0.9181	-0.1010	0.0204	0.0898
	EP	1.0484		0.8218	0.0000	0.0412	0.3888
	AEP	1.1942, 0.9730		0.8256, 0.8353	0.0000	0.0269	0.2098
12 (midnight)	α -Stable	1.6831	0.1040	0.5425	0.0156	0.0258	0.0346
	NIG	0.5396	0.0012	0.6859	-0.0015	0.0216	0.1137
	EP	0.8146		0.6737	0.0000	0.0277	0.3375
	AEP	0.7958, 0.8584		0.6795, 0.6768	0.0000	0.0333	0.3597

Tail: α (stable and NIG), b (EP), b_l and b_r (AEP). *Skewness:* β (stable and NIG). *Scale:* σ (stable and NIG), a (EP), a_l and a_r (AEP). *Location:* μ . Asymptotic 5% limiting values: 0.0356 (D), 0.443 (W^2). Data in bold type indicate the minimum goodness-of-fit test values.

TABLE 25.4 Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions Fitted to the Filtered and Rescaled Log-Returns: Powernext

Auction	Distribution	Tail	Parameter		Location	Test Value	
			Skewness	Scale		D	W^2
4 a.m.	α -Stable	1.8312	-0.4342	0.6038	0.0082	0.0120	0.0510
	NIG	1.2566	-0.2787	1.1539	0.2624	0.0181	0.0828
	EP	1.1797		0.7905	0.0000	0.0480	0.7975
	AEP	1.0566, 1.6008		0.8201, 0.8150	0.0000	0.0288	0.2274
	α -Stable	1.7688	-0.6255	0.5946	-0.0235	0.0137	0.0577
8 a.m.	NIG	1.0771	-0.2594	1.0354	0.2569	0.0108	0.0224
	EP	1.1690		0.7963	0.0000	0.0490	0.9730
	AEP	1.0135, 1.4289		0.8072, 0.7973	0.0000	0.0229	0.1169
	α -Stable	1.7186	-0.2653	0.5860	-0.0295	0.0123	0.2397
	NIG	0.8389	-0.0525	0.9327	0.0585	0.0128	0.0568
12 (noon)	EP	1.0801		0.7820	0.0000	0.0300	0.2743
	AEP	1.0479, 1.1127		0.7836, 0.7808	0.0000	0.0238	0.2124
	α -Stable	1.7744	-0.5740	0.6182	-0.0558	0.0140	0.3900
	NIG	1.0057	-0.1229	1.0977	0.1351	0.0117	0.0513
	EP	1.2018		0.8338	0.0000	0.0354	0.4203
4 p.m.	AEP	1.1266, 1.2595		0.8331, 0.8281	0.0000	0.0261	0.2295
	α -Stable	1.7473	-0.2874	0.6255	-0.0241	0.0141	0.1306
	NIG	0.8779	-0.0660	1.0430	0.0786	0.0098	0.0255
	EP	1.1305		0.8325	0.0000	0.0263	0.2063
	AEP	1.0786, 1.1721		0.8319, 0.8280	0.0000	0.0185	0.1200

(Continued)


TABLE 25.4 (Continued) Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions Fitted to the Filtered and Rescaled Log>Returns: Powernext

Auction	Distribution	Tail	Parameter			Test value	
			Skewness	Scale	Location	D	W^2
12 (midnight)	α -Stable	1.8102	-0.6648	0.6420	-0.0294	0.0122	0.1364
	NIG	1.1885	-0.2262	1.2658	0.2454	0.0138	0.0673
	EP	1.3072		0.8767	0.0000	0.0348	0.5336
	AEP	1.1684, 1.5158		0.8817, 0.8790	0.0000	0.0175	0.1478

Tail: α (stable and NIG), b (EP), b_l and b_r (AEP). *Skewness:* β (stable and NIG). *Scale:* α (stable and NIG), a (EP), a_l and a_r (AEP). *Location:* μ . Asymptotic 5% limiting values: 0.0318 (D), 0.443 (W^2). Data in bold type indicate the minimum goodness-of-fit test values.

estimates of the skewness parameter, the Powernext left tails are longer in all auctions. Note that the point estimates of a_r and a_l are very similar to each other in all markets; as an implication, the skewness in log-returns is mainly due to tail asymmetries, not so much to asymmetries in scale.

The reported goodness-of-fit criteria are evidence in support of the α -stable and NIG models, which outperform the EP and AEP distributions. Both the α -stable and NIG laws provide excellent fits, but neither clearly prevails over the other, as both have some success in a certain number of hourly auctions. In fact, in some hours the NIG fits better according to one goodness-of-fit criterion, while the α -stable fits better according to the other criterion. Note further that the (symmetric) EP distribution virtually always provides the worst fit, presumably because of its inability to capture the skewness. Finally, although the AEP is outperformed by the stable and NIG laws, the goodness-of-fit statistics for the AEP distribution are often below the 5% critical values. The fitting performances of the α -stable, NIG, EP and AEP distributions in a selected hourly auction (4 p.m.) can be seen in Figure 25.3.

**FIGURE 25.3** Density fit of filtered and rescaled log-returns: NordPool, APX and Powernext 4 p.m. auctions.

25.5 Robustness

Unlike stock prices and exchange rates, the electricity price process cannot be approximated by a geometric random walk. Indeed, there is no widespread support in favour of either multiplicative or additive representations of electricity price processes. It is also quite clear that the price dynamics is driven by multiple seasonal factors and is possibly subject to regime shifts. As major implications, there is no unambiguous way to define price returns, neither can one determine beforehand what method is most appropriate to deseasonalize the data. Getting the tails and skewness ‘right’, a major issue in risk management, is not easy, as the estimated distributional parameters may change depending on the definition of returns and on the filtering methodology. These issues are addressed in this section, where the robustness of the results is assessed along two lines: (i) changing the definition of returns while keeping the same deseasonalization method as before (Cholesky filter plus volatility–price rescaling); and (ii) keep working on log-returns while changing the deseasonalization method.

There are convincing reasons to expect that the shape of the returns distribution may be sensitive to the very definition of returns. Log-returns are usually considered as approximations of the percentage returns, defined as $(p_t - p_{t-1})/p_{t-1}$. The difference between a log-return x_t and a percentage return is of the order of $\frac{1}{2}x_t^2 + \frac{1}{6}x_t^3 + \dots$ (Eberlein and Keller 1995). While this difference is negligible in financial markets, the approximation may be quite poor in power exchanges, due to the extremely large magnitudes of electricity price fluctuations. A second problem is that the logarithmic transformation dampens the extreme returns and makes the distribution of returns more symmetric, thereby affecting the estimated shape and skewness parameters. Empirically, this effect has been verified by Weron (2009), who finds that the distribution of price changes $p_t - p_{t-1}$ displays heavier tails than the distribution of log-returns. For the above reasons, the robustness of the foregoing results needs to be checked using alternative definitions of price returns.

Table 25.5 reports the estimation results when the α -stable, NIG, EP and AEP models are fitted on the empirical probability densities of percentage returns and price changes, Cholesky-filtered and rescaled as in Section 25.2.* The first result is that, in most hours, the distributions of price changes and percentage returns have heavier tails than the distributions of logarithmic returns: indeed, the point estimates of the shape parameters are lower. There are some exceptions concerning APX and Powernext: in those markets, the percentage returns are less fat-tailed than the log-returns from 4 a.m. to 4 p.m., and the tails of the price changes decay faster during the night. Second, the skewness patterns observed for log-returns are not robust when one considers percentage returns and price changes. For instance, the NordPool and Powernext price change distributions are closer to symmetric, but there is a negative skew in the APX (APX log-returns, on the contrary, display a positive skew). Less clear are the results for percentage returns, with alternating signs, although the magnitudes of the skewness parameters tend to be mild. Different distributions give contrasting results: the Powernext AEP left shape parameter b_l is now often slightly larger than the right shape parameter b_r , hinting at a positive skew, but α -stable skewness parameters for the same market are often negative. These findings are only partly in line with the evidence of Weron (2009), and show that analysing returns of individual hours can shed light on intricate intraday patterns. Finally, the comparative performance of the α -stable law improves for the percentage returns and for the NordPool price changes (i.e. it is the best-fitting distribution in a larger number of hourly auctions), while the NIG prevails as a description of APX and Powernext price changes. The AEP law is now the best-fitting distribution in two instances (APX 4 a.m. percentage returns, Powernext 4 a.m. price changes).

The results presented in Section 25.4 may also be sensitive to the filtering methodology. On the one hand, if power-law volatility–price scaling is an effective way of controlling for heteroskedasticity, the time series of log-returns filtered using only the Cholesky algorithm should be the superposition of underlying homoskedastic (and possibly non-heavy tailed) series. It should therefore display longer

* Information on the scale and location parameters, which is less essential, is omitted in order to save space.

TABLE 25.5 Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions. Percentage Returns: Cholesky-Filtered and Rescaled Percentage Returns

Auction	Distribution	NordPool				APX				Powernext			
		Tail	Skewness	D	W^2	Tail	Skewness	D	W^2	Tail	Skewness	D	W^2
Percentage returns													
4 a.m.	α -Stable	1.7601	0.0357	0.0082	0.0207	1.3853	−0.7461	0.0477	0.5850	1.6046	−0.5874	0.0229	0.1364
	NIG	0.9289	0.0098	0.0119	0.0526	0.1848	0.0043	0.0473	1.0026	0.5143	0.0694	0.0306	0.2866
	EP	1.1555		0.0184	0.1873	0.5968		0.0535	1.2675	0.7526		0.0454	0.9482
	AEP	1.1667, 1.1465		0.0178	0.1859	0.8469, 1.7175		0.0469	0.6978	0.7423, 0.7555		0.0486	0.9994
8 a.m.	α -Stable	1.6227	−0.0139	0.0176	0.1590	1.2095	−0.1256	0.0221	0.1497	1.6464	0.1672	0.0209	0.1280
	NIG	0.5102	0.0523	0.0166	0.0713	0.0425	0.0122	0.0388	0.6097	0.5762	0.0147	0.0199	0.2110
	EP	0.8235		0.0411	0.7424	0.4581	0.0000	0.1233	6.8636	0.8996		0.0349	0.5283
	AEP	0.9461, 0.7715		0.0296	0.2735	0.7217, 1.0243		0.0431	0.8278	0.9089, 0.8938		0.0326	0.5096
12 (noon)	α -Stable	1.7699	0.1106	0.0147	0.0624	1.1537	0.2139	0.0171	0.0672	1.4329	−0.1575	0.0139	0.0766
	NIG	0.7739	0.1199	0.0191	0.1656	0.1066	0.0229	0.0159	0.0535	0.4352	−0.0256	0.0218	0.1382
	EP	0.9735		0.0408	0.9430	0.5422		0.0894	2.5538	0.7390		0.0329	0.4407
	AEP	1.2500, 0.9019		0.0283	0.4420	0.5328, 0.5818		0.0384	0.5274	0.7288, 0.7464		0.0314	0.4142
4 p.m.	α -Stable	1.7506	0.1628	0.0139	0.0886	1.2728	0.1556	0.0162	0.0568	1.5934	−0.1482	0.0127	0.0651
	NIG	0.7985	0.1359	0.0143	0.0403	0.1701	0.0462	0.0235	0.1966	0.4663	0.0021	0.0199	0.1148
	EP	1.0205		0.0430	0.9258	0.5650		0.0996	3.7912	0.7762		0.0286	0.4574
	AEP	1.2837, 0.9251		0.0228	0.2031	0.6743, 0.5689		0.0436	0.7330	0.8105, 0.7597		0.0330	0.4611
8 p.m.	α -Stable	1.7344	0.0949	0.0120	0.0909	1.6523	0.3390	0.0185	0.0585	1.8009	0.2397	0.0113	0.0235
	NIG	0.7287	0.0928	0.0137	0.0728	0.7538	0.2542	0.0274	0.1576	0.9521	0.1711	0.0165	0.0958
	EP	1.0007		0.0376	0.5930	0.9065		0.0781	2.4277	1.0947		0.0385	0.6816
	AEP	1.1944, 0.9284		0.0232	0.2478	1.6815, 0.8225		0.0299	0.2811	1.3845, 0.9950		0.0255	0.2554
12 (midnight)	α -Stable	1.8025	−0.2344	0.0157	0.1124	1.7001	−1.0000	0.1330	6.3918	1.8899	−0.2964	0.0149	0.0334
	NIG	1.0913	−0.0772	0.0128	0.0392	0.4942	0.1024	0.0221	0.1155	1.4407	−0.0210	0.0113	0.0285
	EP	1.2521		0.0245	0.1600	0.4621		0.4084	99.4721	1.4470		0.0155	0.0711
	AEP	1.1997, 1.3100		0.0197	0.0674	0.6949, 3.3029		0.1279	3.5804	1.4595, 1.4376		0.0168	0.0793
Price changes													
4 a.m.	α -Stable	1.6971	−0.2596	0.0183	0.1204	1.7924	−0.3495	0.0162	0.0487	1.8861	0.7159	0.0172	0.1173
	NIG	0.7944	−0.0678	0.0132	0.0872	1.0989	−0.1570	0.0117	0.0263	1.6595	0.2079	0.0152	0.0807
	EP	0.6237		0.0322	0.4607	1.2815		0.0266	0.2532	1.4724		0.0163	0.1359
	AEP	0.9799, 1.0960		0.0221	0.3037	1.0236, 1.1981		0.0153	0.0616	1.6274, 1.3686		0.0190	0.0680
8 a.m.	α -Stable	1.6136	−0.0794	0.0124	0.0397	1.7186	−0.2948	0.0374	0.1449	1.7706	−0.2979	0.0314	0.1547
	NIG	0.4924	0.0186	0.0208	0.2354	0.7468	−0.0124	0.0224	0.1915	0.9708	−0.0093	0.0110	0.0302
	EP	0.7959		0.0375	0.8018	0.9950		0.0352	0.4555	1.1434		0.0180	0.1221

(Continued)

TABLE 25.5 (Continued) Parameter Estimates and Goodness-of-Fit Statistics for α -Stable, NIG, EP and AEP Distributions. Percentage Returns: Cholesky-Filtered and Rescaled Percentage Returns

Auction	Distribution	NordPool				APX				Pownernext			
		Tail	Skewness	D	W^2	Tail	Skewness	D	W^2	Tail	Skewness	D	W^2
12 (noon)	AEP	0.8523, 0.7687		0.0379	0.6920	0.1031, 0.3799		0.0417	0.4929	1.1850, 1.1193		0.0209	0.1504
	α -Stable	1.7239	0.1265	0.0167	0.0891	1.3951	−0.0885	0.0368	0.1598	1.5186	−0.0255	0.0324	0.0892
	NIG	0.6638	0.0758	0.0191	0.1708	0.3059	0.0743	0.0239	0.1747	0.4351	0.0440	0.0231	0.1789
	EP	0.8814		0.0469	0.9318	0.6739		0.0781	2.6601	0.7208		0.0494	1.0317
	AEP	1.0399, 0.8474		0.0352	0.5409	0.4287, 0.1061		0.0289	0.2943	0.8093, 0.6770		0.0319	0.4685
4 p.m.	α -Stable	1.7237	0.1361	0.0278	0.1118	1.3386	−0.1266	0.0307	0.2424	1.5252	−0.1906	0.0453	0.8678
	NIG	0.7389	0.1173	0.0169	0.1295	0.2013	0.0478	0.0300	0.1919	0.3569	0.0079	0.0311	0.2986
	EP	0.2201		0.0464	1.0360	0.5801		0.0953	3.3317	0.7068		0.0392	0.7774
	AEP	1.1818, 0.8735	0.0752	0.0320	0.3599	0.4019, 0.1754		0.0336	0.4289	0.7435, 0.6884		0.0366	0.6522
8 p.m.	α -Stable	1.7509	0.0400	0.0320	0.1361	1.1560	−0.2662	0.0176	0.0514	1.7136	−0.0556	0.0143	0.0701
	NIG	0.6660		0.0189	0.1407	0.1504	0.0283	0.0199	0.1369	0.7078	0.0730	0.0201	0.1237
	EP	0.4473		0.0326	0.4684	0.5909		0.0789	2.2140	0.9454		0.0420	0.6012
	AEP	0.9859, 0.8856		0.0234	0.3973	0.9745, 0.8882		0.0308	0.2658	1.0777, 0.8835		0.0272	0.3693
12 (midnight)	α -Stable	1.7473	−0.1336	0.0299	0.1708	1.6631	0.2094	0.0398	0.1232	1.8196	−0.1083	0.0310	0.1257
	NIG	0.9149	−0.0403	0.0087	0.0261	0.8023	0.2030	0.0187	0.1034	1.1197	0.0230	0.0114	0.0388
	EP	0.4005		0.0192	0.1206	0.9733		0.0596	1.2475	1.2703		0.0172	0.0776
	AEP	1.1200, 1.1702		0.0161	0.0817	0.0764, 0.2028		0.0240	0.2088	1.3464, 1.2336		0.0181	0.1131

Tail: α (stable and NIG), b (EP), b_l and b_r (AEP). Skewness: β (stable and NIG). Asymptotic 5% limiting values: 0.0291 (D , NordPool), 0.0356 (D , APX), 0.0318 (D , Pownernext), 0.443 (W^2). Data in bold type indicate the minimum goodness-of-fit test values.

tails. One can therefore assess the impact of heteroskedasticity on the distributional shapes by fitting the empirical density functions of the log-returns after applying only the Cholesky filter (i.e. x), without considering volatility–price scaling. On the other hand, the Cholesky filter is built upon the sample autocovariance function, hence it may suffer from the same limitations as Fourier-based filters if the time series is non-stationary. The outcome of applying a filter that is less suited to deal with non-stationarities would probably be a mixture of random variables, leading to spurious estimates of the asymmetry and tail parameters. Wavelet filters should be immune to these problems. Following Weron (2006, 2009), one approximates the long-term seasonal component of the time series of electricity prices by means of an S8 approximation, based on a Daubechies-20 wavelet filter; then the long-term seasonal component is subtracted from the time series and the weekly pattern is removed using a moving-average filter. The minimum of the resulting time series is aligned with the minimum of the original price series, and then the log-returns are computed, ready for the α -stable, NIG, EP and AEP fit.

The estimates in Table 25.6 refer to Cholesky-filtered and wavelet-filtered log-returns. These estimates show that, compared with the distributions of Cholesky-filtered and rescaled log-returns, the tails are at least as fat. One noteworthy exception is the NIG distribution, whose shape parameter estimates are unusually high. The skewness estimates in the case of Cholesky-filtered log-returns are basically confirmed for APX (positive) and Powernext (negative), although the magnitudes are milder; NordPool log-returns are now approximately symmetric. For the wavelet-filtered log-returns, the skewness is now positive in the NordPool (it was often negative for Cholesky-filtered and rescaled log-returns), mildly negative in APX (it was positive), and the negative sign in the Powernext is confirmed in most hours, but the magnitudes are smaller. Finally, the α -stable law now displays an improved fitting performance and dominates the other distributions in a greater number of hours. The AEP provides the best fit twice (4 a.m. and 8 p.m.) in the Powernext market.

25.6 Conclusion

This paper contributes to the characterization of the probability density of the price returns in some European day-ahead electricity markets (NordPool, APX, Powernext) by fitting flexible and general families of distributions, such as the α -stable, Normal Inverse Gaussian (NIG), Exponential Power (EP) and Asymmetric Exponential Power (AEP), and comparing their goodness of fit. One finds that the probability of observing extremely large (positive or negative) returns is higher than for the Gaussian, confirming a very robust finding in the literature. The goodness-of-fit tests suggest that the α -stable and the NIG systematically outperform the EP and AEP models, although no clear ranking can be established between the two ‘winners’. The evidence of heavy tails is robust to changing the definition of returns (from log-returns to percentage returns and price changes) and the deseasonalization methodology. Yet, the point estimates differ across cases; in particular, the logarithmic transform and volatility rescaling tend to dampen the extreme returns. Both the skewness and the tail behaviour vary across markets, and one also observes interesting intra-daily patterns which were not detected in previous studies that focused on average daily returns.

The evidence of heavy tails is in accordance with the intuition behind the regime-switching and jump-diffusion models. If the returns can be represented as scale–location mixtures—as the good performance of the NIG distribution seems to suggest—the market is characterized by low volatility most of the time, but instances of extreme volatility are not negligibly rare. This is very much consistent with the idea that power exchanges undergo transitions between quiet and turbulent price regimes. The better fitting performance of the α -stable and NIG laws compared with the EP and AEP models suggests that there is a time scaling between daily and longer-term market risk: both the α -stable and the NIG distributions are closed under convolution, therefore their shapes are preserved under time aggregation. Moreover, to the extent that the α -stable distribution outperforms the NIG, the results signal the non-convergence of the second moment. This is bad news vis-à-vis the use of price volatility as a measure of price risk in power exchanges. The results also suggest that the skewness and the

TABLE 25.6 Parameter Estimates and G-of-Fit Statistics for $-\alpha$ -Stable, NIG, EP and AEP Distributions. *Cholesky-Filtered*: Cholesky-Filtered Log-Returns. *Wavelet-Filtered*: Wavelet-Filtered Log-Returns

Auction	Distribution	NordPool				APX				Powernext			
		Tail	Skewness	D	W ²	Tail	Skewness	D	W ²	Tail	Skewness	D	W ²
Cholesky-filtered													
4 a.m.	α-Stable	1.5964	0.0079	0.0167	0.0860	1.5781	0.2587	0.0284	0.2567	1.8050	−0.4188	0.0125	0.0435
	NIG	0.7127	−0.0578	0.0230	0.2309	0.7232	−0.1685	0.0396	0.3625	1.1979	−0.2509	0.0164	0.0904
	EP	0.9215		0.0355	0.7323	0.8370		0.0666	1.9110	1.1230		0.0491	0.9247
	AEP	0.8728, 0.9904		0.0304	0.5458	0.7553, 1.2110		0.0445	0.5167	1.0020, 1.4278		0.0259	0.2578
8 a.m.	α-Stable	1.6432	0.0686	0.0147	0.0549	1.5114	0.3584	0.0355	0.4784	1.7001	−0.4300	0.0119	0.0496
	NIG	0.7614	0.0052	0.0133	0.0575	0.5966	−0.1339	0.0348	0.3671	0.9606	−0.2285	0.0136	0.0538
	EP	0.9515		0.0213	0.2579	0.7781		0.0697	2.1414	1.0416		0.0483	1.1732
	AEP	0.9602, 0.9422		0.0232	0.2574	0.6994, 1.1055		0.0399	0.4188	0.9146, 1.3346		0.0190	0.1437
12 (noon)	α-Stable	1.7258	0.0659	0.0103	0.0351	1.7843	0.6160	0.0194	0.0802	1.6686	−0.1333	0.0101	0.0229
	NIG	0.9168	0.0559	0.0129	0.0577	1.2039	0.2783	0.0137	0.0418	0.8066	−0.0267	0.0120	0.0532
	EP	1.0639		0.0268	0.2820	1.2328		0.0466	0.5844	1.0165		0.0257	0.2312
	AEP	1.1297, 1.0181		0.0187	0.2149	1.5485, 1.0719		0.0150	0.0544	1.0065, 1.0216		0.0242	0.2242
4 p.m.	α-Stable	1.6814	0.0106	0.0107	0.0363	1.6672	0.1706	0.0191	0.0519	1.6853	−0.3301	0.0111	0.0240
	NIG	0.8287	0.0436	0.0153	0.1184	0.8550	0.1176	0.0184	0.0775	0.8761	−0.1011	0.0126	0.0542
	EP	1.0130		0.0272	0.4278	1.0258		0.0411	0.4397	1.0752		0.0352	0.4411
	AEP	1.0499, 0.9746		0.0221	0.3581	1.1896, 0.9489		0.0216	0.1692	1.0056, 1.1325		0.0238	0.2382
8 p.m.	α-Stable	1.6941	0.0087	0.0091	0.0186	1.6988	0.1062	0.0148	0.0604	1.6705	−0.0576	0.0118	0.0530
	NIG	0.8317	0.0280	0.0150	0.1060	0.9226	0.0907	0.0193	0.0950	0.8416	−0.0813	0.0147	0.0516
	EP	0.9868		0.0260	0.3963	1.1021		0.0393	0.3406	1.0505		0.0285	0.3461
	AEP	1.0166, 0.9593		0.0251	0.3643	1.2435, 1.0287		0.0256	0.2134	0.9805, 1.1437		0.0233	0.1708
12 (midnight)	α-Stable	1.6232	−0.0442	0.0099	0.0285	1.7510	0.2557	0.0232	0.0718	1.7033	−0.3023	0.0126	0.0531
	NIG	0.7315	−0.0469	0.0133	0.0627	0.9172	0.0250	0.0227	0.1257	0.9383	−0.1440	0.0160	0.0831
	EP	0.9572		0.0296	0.3633	0.8640		0.0305	0.3945	1.1021		0.0408	0.5833
	AEP	0.9181, 1.0044		0.0208	0.2555	0.8475, 0.9016		0.0357	0.4419	1.0050, 1.2632		0.0206	0.2059
Wavelet-filtered													
4 a.m.	α-Stable	1.3899	0.1277	0.0193	0.074	1.6652	0.1434	0.0124	0.0224	1.6328	0.0079	0.0253	0.1164
	NIG	9.2782	0.0463	0.0201	0.0904	3.3436	0.1840	0.0272	0.1785	4.9146	0.0844	0.0152	0.0249
	EP	0.7769		0.0312	0.2811	0.6982		0.0498	0.5507	1.0240		0.0191	0.0339

(Continued)

TABLE 25.6 (Continued) Parameter Estimates and G-of-Fit Statistics for $-\alpha$ -Stable, NIG, EP and AEP Distributions. *Cholesky-Filtered*: Cholesky-Filtered Log>Returns. *Wavelet-Filtered*: Wavelet-Filtered Log>Returns

Auction	Distribution	NordPool				APX				Powernext			
		Tail	Skewness	D	W^2	Tail	Skewness	D	W^2	Tail	Skewness	D	W^2
8 a.m.	AEP	0.7640, 0.7859		0.0305	0.2781	0.6850, 0.7121		0.0442	0.5128	1.0026, 1.0458		0.0147	0.0304
	α -Stable	1.3670	0.0655	0.0230	0.0638	1.6838	0.1257	0.0253	0.1012	1.7980	0.1765	0.0237	0.0575
	NIG	5.7047	0.3095	0.0203	0.0873	3.4372	0.2307	0.0375	0.4312	7.8396	0.4915	0.0274	0.1315
	EP	0.7312		0.0347	0.3582	0.7567		0.0576	0.9541	1.2164		0.0352	0.2467
12 (noon)	AEP	0.7485, 0.7111		0.0300	0.3124	0.7379, 0.7779		0.0511	0.9002	1.1703, 1.2707		0.0308	0.226
	α -Stable	1.5447	0.2860	0.0200	0.0515	1.0698	0.1326	0.0241	0.1144	1.5319	0.0961	0.0238	0.0880
	NIG	11.6727	0.8902	0.0199	0.0584	1.6736	0.0037	0.0168	0.0372	3.9648	0.1671	0.0319	0.29
	EP	0.901		0.0393	0.2715	0.5230		0.0298	0.2460	0.7398		0.0467	0.6475
4 p.m.	AEP	0.9381, 0.8586		0.0295	0.1976	0.5221, 0.5238		0.0291	0.2458	0.7200, 0.7638		0.0420	0.6102
	α -Stable	1.5365	0.1359	0.0167	0.0329	1.0541	0.0543	0.0268	0.1109	1.5584	0.2774	0.0208	0.0863
	NIG	12.9554	1.5886	0.0201	0.0705	0.5572	0.0024	0.0215	0.1004	5.3607	0.0180	0.0374	0.296
	EP	0.8617		0.0474	0.4613	0.4238		0.0462	0.6175	0.7932		0.0462	0.6001
8 p.m.	AEP	0.9379, 0.7957		0.0273	0.2256	0.4355, 0.4120		0.0473	0.5795	0.7867, 0.7990		0.0460	0.5902
	α -Stable	1.6042	0.2205	0.0177	0.0345	1.2354	0.0548	0.0244	0.1109	1.7631	0.3230	0.0207	0.0440
	NIG	16.8065	0.0877	0.0341	0.1956	2.7449	0.0124	0.0331	0.1757	12.0221	0.0761	0.0295	0.2032
	EP	0.8731		0.0455	0.3815	0.6965		0.0459	0.3761	1.1684		0.0347	0.2634
12 (midnight)	AEP	0.9081, 0.8376		0.0359	0.3130	0.6622, 0.7336		0.0295	0.2206	1.0527, 1.3768		0.0184	0.0562
	α -Stable	1.5117	0.2004	0.0196	0.0673	1.4497	0.1648	0.0274	0.0868	1.7528	0.3078	0.0247	0.0734
	NIG	15.4975	0.4670	0.0261	0.1362	2.3668	0.0098	0.0321	0.2113	21.0111	2.3941	0.0180	0.0359
	EP	0.8135		0.0342	0.3425	0.6598		0.0468	0.5054	1.1671		0.0330	0.2026
	AEP	0.8190, 0.8061		0.0341	0.3371	0.6696, 0.6480		0.0410	0.4851	1.0810, 1.2551		0.0184	0.0649

Tail: α (stable and NIG), b (EP), b_l and b_r (AEP). Skewness: β (stable and NIG). Asymptotic 5% limiting values: 0.0291 (D , NordPool), 0.0356 (D , APX), 0.0318 (D , Powernext), 0.443 (W^2). Data in bold type indicate the minimum goodness-of-fit test values

kurtosis are at least as important for risk management as volatility, in line with the theoretical results obtained by Bessembinder and Lemmon (2002) concerning the pricing of power forwards. Still, in interpreting the results one has to acknowledge the relative scarcity of data for the tails of electricity returns distributions, compared with the wealth of high-frequency financial data. This poses estimation and goodness-of-fit problems, as pointed out by Weron *et al.* (2004). Such problems are going to persist unless we wait long enough until many years of data are available. Thus, estimation on data from power exchanges established more recently is not expected to yield better results, at least not in the near future.

The results reported in this paper can be seen as the starting point for further work. One could extend the policy-oriented analysis performed by Robinson and Baniak (2002) on the impact of Contracts for Differences (CfDs) and test the effects of further policy measures, such as the introduction of the EU ETS scheme for carbon emissions and the liberalization of retail trading.

References

- Andrews, D.F. and Mallows, C.L., Scale mixtures of Normal distributions. *J. R. Statist. Soc., Ser. B*, 1974, **36**(1), 99–102.
- Barndorff-Nielsen, O.E., Normal Inverse Gaussian distributions and stochastic volatility modelling. *Scand. J. Statist.*, 1997, **24**(1), 1–13.
- Bellini, F., Empirical analysis of electricity spot prices in European deregulated markets, *Quaderni Ref.* 7/2002, 2002.
- Bessembinder, H. and Lemmon, M.L., Equilibrium pricing and optimal hedging in electricity forward markets. *J. Finance*, 2002, **57**(3), 1347–1382.
- Borak, S., Härdle, W. and Weron, R., Stable distributions. SFB 649 Discussion Paper 2005–008, Humboldt Universität zu Berlin, 2005.
- Bosco, F., Parisio, L. and Pelagatti, M., Deregulated wholesale electricity prices in Italy: An empirical analysis. *Int. Adv. Econ. Res.*, 2007, **13**, 415–432.
- Bottazzi, G., Subbotools: A reference manual. LEM Working Paper 2004/14, S. Anna School of Advanced Studies, Pisa, 2004.
- Bottazzi, G. and Sapio, S., Power exponential price returns in day-ahead power exchanges. In *Econophysics and Sociophysics of Markets and Networks*, edited by A. Chatterjee and B.K. Chakrabarti, 2007 (Springer: Berlin).
- Bottazzi, G., Sapio, S. and Secchi, A., Some statistical investigations on the nature and dynamics of electricity prices. *Physica A*, 2005, **355**(1), 54–61.
- Bottazzi, G. and Secchi, A., Why are distributions of firm growth rates tent-shaped? *Econ. Lett.*, 2003, **80**, 415–420.
- Bottazzi, G. and Secchi, A., Maximum Likelihood estimation of the symmetric and Asymmetric Exponential Power distribution. LEM, S. Anna School of Advanced Studies, Pisa, 2007.
- Byström, H., Extreme value theory and extremely large electricity price changes. *Int. Rev. Econ. Finance*, 2005, **14**(1), 41–55.
- Chan, K.F. and Gray, P., Using extreme value theory to measure value-at-risk for daily electricity spot prices. *Int. J. Forecast.*, 2006, **22**, 283–300.
- Choy, S.T.B. and Walker, S.G., The extended exponential power distribution and Bayesian robustness. *Statist. Probab. Lett.*, 2003, **65**, 227–232.
- D'Agostino, R.B. and Stephens, M.A., *Goodness-of-fit Techniques*, 1986 (CRC Press: Boca Raton, FL).
- De Jong, C., The nature of power spikes: A regime-switch approach. *Stud. Nonlinear Dynam. Econometr.*, 2006, **10**(3), (article 3).
- Deng, S.-J. and Jiang, W., Levy process-driven mean-reverting electricity price model: The marginal distribution analysis. *Decis. Supp. Syst.*, 2005, **40**(3/4), 483–494.

- Deng, S.-J., Jiang, W. and Xia, Z., Alternative statistical specifications of commodity price distribution with fat tails. *Adv. Model. Optimiz.*, 2002, 4, 1–8.
- Diebold, F.X., Ohanian, L.E. and Berkovitz, J., Dynamic equilibrium economies: A framework for comparing models and data. Working Paper No. 97-7, Federal Reserve Bank of Philadelphia, 1997.
- Eberlein, E. and Keller, U., Hyperbolic distributions in finance. *Bernoulli*, 1995, 1(3), 281–299.
- Eberlein, E. and Stahl, G., Both sides of the fence: A statistical and regulatory view of electricity risk. *Energy Power Risk Mgmt*, 2003, 8, 34–38.
- Fu, D., Pammolli, F., Buldyrev, S., Riccaboni, M., Matia, K., Yamasaki, K. and Stanley, H.E., The growth of business firms: Theoretical framework and empirical evidence. *Proc. Natn Acad. Sci.*, 2005, **102**, 18801–18806.
- Karakatsani, N.V. and Bunn, D., Modelling stochastic volatility in high-frequency spot electricity prices. Department of Decision Sciences, London Business School, 2004.
- Kogon, S.M. and Williams, D.B., Characteristic function based estimation of stable distribution parameters. In *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, edited by R.J. Adler, R.E. Feldman, and M. Taqqu, 1998 (Springer: Berlin).
- Koutrouvelis, I.A., Regression-type estimation of the parameters of stable laws. *J. Am. Statist. Assoc.*, 1980, 75(372), 918–928.
- Mandelbrot, B., The variation of certain speculative prices. *J. Business*, 1963, **34**(4), 394–419.
- Mugele, C., Rachev, S.T. and Trück, S., Stable modeling of different European power markets. *Invest. Mgmt Financial Innov.*, 2005, **3**, 65–85.
- Robinson, T. and Baniak, A., The volatility of prices in the English and Welsh electricity pool. *Appl. Econ.*, 2002, **34**, 1487–1495.
- Samorodnitsky, G. and Taqqu, M.S., *Stable NonGaussian Random Processes*, 1994 (Chapman & Hall: London).
- Scalas, E. and Kim, K., The art of fitting financial time series with Levy stable distributions. *Korean J. Phys.*, 2007, **50**, 105–111.
- Simonsen, I., Volatility of power markets. *Physica A*, 2005, **355**, 10–20.
- Stephens, M.A., EDF statistics for goodness of fit and some comparisons. *J. Am. Statist. Assoc.*, 1974, **69**(347), 730–737.
- Subbotin, M.F., On the law of frequency of errors. *Matematicheskii Sbornik*, 1923, **31**, 296–301.
- Weron, R., Computationally intensive value at risk calculations. In *Handbook of Computational Statistics*, edited by J.E. Gentle, W. Härdle, and Y. Mori, pp. 911–950, 2004 (Springer: Berlin).
- Weron, R., Heavy tails and electricity prices. In *The Deutsche Bundesbank's 2005 Annual Fall Conference on Heavy Tails and Stable Paretian Distributions in Finance and Macroeconomics*, Eltville, 10–12 November 2005.
- Weron, R., *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*, 2006 (Wiley: Chichester).
- Weron, R., Heavy-tails and regime-switching in electricity prices. *Math. Meth. Oper. Res.*, 2009, **69**, 457–473.
- Weron, R., Bierbrauer, M. and Trück, S., Modeling electricity prices: Jump diffusion and regime switching. *Physica A*, 2004, **336**, 39–48.
- West, M., On scale mixtures of Normal distributions. *Biometrika*, 1987, **74**, 646–648.

26

Stochastic Spot Price Multi-Period Model and Option Valuation for Electrical Markets

26.1	Introduction.....	571
26.2	Spot Price Multi-Period Model.....	572
26.3	Model Calibration.....	573
	The Price Forward Curve • The Short-Term Process • The Long-Term Process • The Term Structure of Volatility	
26.4	Scenario Generation and Valuation of Structured Products	578
	Estimation of the Volatility Surface	
26.5	Structured Product Valuation	580
	Valuation and Risk Analysis: TS-Energy • Virtual Power Plant Valuation	
26.6	Conclusion.....	583
	Acknowledgements.....	584
	References.....	584

Eivind Helland

Timur Aka

Eric Winnington

26.1 Introduction

The deregulation of electricity markets has resulted in competitive prices, but also in higher fluctuations of electricity price development since electricity is scarcely storable. Most markets exhibit high price volatility and intermittent price spikes. However, the prices have a strong daily, weekly and yearly periodicity that is explained by looking at the market price as an equilibrium price based on supply and demand curves. Since the demand is very inelastic, the marginal costs of the supply side determine the price to a large extent. If the total load is low, plants with the lowest variable production costs are used, and if the total load is high, gas- or oil-fired plants with high fuel costs are run additionally. The periodicity of the total load is responsible for the periodicity of the electricity prices. The total load has a random component, depending on short-term weather conditions and other uncertain parameters, but it also has a clearly predictable part, and so do electricity prices (Müller *et al.* 2004). This has led to a demand for derivative products that hedge the owner against high prices. A thorough understanding of the stochastic price dynamics is thus necessary for the purposes of risk management and derivative pricing. At Axpo AG, the largest energy company in Switzerland, we have developed a stochastic spot price multi-period model for optimising hydro power plant production, pricing complex options and structured products, and hedging the asset portfolio.

A stochastic spot price model has been developed in order to evaluate standard and off-standard products in the electrical market and for the optimisation of power production in pumped storage hydro parks. The spot price model is built up in several steps and calibration is carried out based on the price distribution curves of spot prices and the log returns of forward contracts as well as standard option contracts over several delivery periods.

The base price level of the stochastic spot price model is an hourly price forward curve that has been modelled using a regression model based on past European Exchange (EEX) price data. The main input parameters were the calendar dates (weekdays, holidays, etc.) to obtain seasonality effects (intraday, weekly and yearly), historical meteorological data (daily temperature, wind and precipitation data) and long-term market expectations. The price forward curve (PFC) was then scaled such that the average price of the hourly PFC was equal to the future power contract at the corresponding time interval to obtain the arbitrage-free condition. Secondly, the daily changes of the PFC have been modelled based on a log-normal, mean-reverting (weekly) process of the front-year forward base contract in order to reflect the shift in spot prices due to changes in long-term forward contracts. Thirdly, a lognormal, mean-reverting, short-term (hourly) process was modelled to represent the intra-day and day-ahead perturbations in spot prices. Additionally, two jump process are admixed to model the positive and negative jumps observed on the German market.

Based on the most recent PFC, and a calibrated set of weekly, hourly and jump processes, a set of price scenarios is generated using the spot simulation model. These scenarios are used to evaluate the price of future options contracts available on the EEX. These contracts are created and evaluated in TS-Energy (software developed by Time-Steps). The distribution of the scenarios' end values are used to evaluate the fair price of a call and a put option based on selected strike prices. These values are compared with the observed prices for call options on the EEX or broker data available for the German market. The stochastic spot price multi-period models are calibrated so that calculated option prices for monthly, quarterly and annual contracts fall within the quoted bid–ask spread.

26.2 Spot Price Multi-Period Model

Our spot price simulation model captures the typical features observed in the electricity market as described above in the Introduction. In this paper we calibrate it to data for German electricity prices on the European Energy Exchange (EEX), but it can readily be adapted to any other market. Therefore, we describe the spot market price by a discrete-time stochastic process with an hourly granularity. The full model can be considered as a three-factor model in order to fully capture the different features observed in electricity markets.

- The base price level of the spot price model is an hourly price forward curve that has been modelled using a regression model accounting for the seasonal patterns and periodicities and scaled with liquid future power contracts.
- A short-term process with positive and negative jumps is used to represent the intraday and day-ahead perturbations in spot prices which account for unexpected variations in supply and demand due to weather conditions and production outages.
- Along-term process including a volatility term structure, which accounts for the uncertainty in supply and demand, fuel costs and macroeconomic variables in the long term.

The fundamental equation of our model can be written as

$$S_t = (X_t + L_{p,t} + L_{op,t}) \frac{Y_t}{Y_0}. \quad (26.1)$$

The factors X_p , $L_{p,t}$ and $L_{op,t}$ produce the short-term variations in price behaviour, whereas the factor Y_t is responsible for its long-term variation. More precisely, the process X_t describes the deviation of spot prices from the hourly price forward curve θ_p , which reflects the current market situation, and is given by

$$X_t = \theta_t \left(1 + \eta \cdot \frac{X_{t-1} - \theta_{t-1}}{\theta_{t-1}} + \sigma_{X,t} \cdot \varepsilon_{X,t} + a \cdot \ln \left(\frac{\theta_{t+24}}{\theta_t} \right) \Delta t_h \right), \quad (26.2)$$

where ε_t are standard normally distributed random numbers, i.e. $\varepsilon_{X,t} \sim \mathcal{N}(0,1)$. η is the autocorrelation parameter of the previous hourly shift of the intraday spot price versus the price forward curve. If the spot price is higher than the expected price forward curve level then higher prices will also be experienced in the following hours. The jump processes $L_{k,t}$, $k \in \{p, op\}$ take into account occasionally observed spikes in the hourly spot prices. Introducing two independent jump processes allows us to properly distinguish the jump intensities in peak and off-peak hours. The discrete jump processes have the form

$$L_k = \mathbf{1}_{\{u_{k,t} < \lambda_{k,t}\}} \cdot \varepsilon_{k,t} \cdot h_k \cdot \mathbf{1}_{\{t \in \mathcal{T}_k\}} + L_{k,t-1} \cdot d_k, \quad k \in \{p, op\}, \quad (26.3)$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the usual indicator function, the disjoint sets $\mathcal{T}_p$ and $\mathcal{T}_{op}$ subdivide hourly time grid nodes into peak (p) and off-peak (op) hours, $u_{k,t}$ are uniformly distributed random numbers, i.e. $u_{k,t} \sim \mathcal{U}(0,1)$ and $\varepsilon_{k,t} \sim \mathcal{N}(0,1)$. The volatility $\sigma_{X,t}$ in Equation (26.2) and the jump intensities $\lambda_{k,t}$ in (26.3) are allowed to be time dependent in order to consider seasonality effects. Finally, the evolution of the long-term variation process is given by

$$Y_{t+1} = \begin{cases} Y_t \cdot \left(1 + 1/7 \left(\ln \left(\frac{K}{Y_t} \right) \Delta t_d + \sigma_{Y,t} \varepsilon_{Y,t} \right) \right), & \text{if } t \in \{0, 7, 14, 21, 28, \dots\} \\ Y_t + \Delta Y_t, & \text{else,} \end{cases} \quad (26.4)$$

where K is a user-defined long-term reference price. Note that the time-step length Δt_d of one day in (26.4) differs from Δt_h in (26.2), which is based on an hourly time grid. Therefore, the long-term process exhibits a random change in its value once a week based on $\varepsilon_{Y,t} \sim \mathcal{N}(0,1)$ and is updated by a constant value $\Delta Y_t := Y_t - Y_{t-1}$ in the remaining time steps. Again, the daily volatility $\sigma_{Y,t}$ in Equation (26.4) is allowed to be a function of time. We point out that all random variables in the above equations are assumed to be stochastically independent.

26.3 Model Calibration

The model is calibrated separately for each factor. The short-term and long-term factor calibrations are described in detail below. For the statistical analysis in the following sections, we used EEX prices from January 1, 2011 to January 1, 2012.

26.3.1 The Price Forward Curve

The price forward curve (PFC) reflects the expectation of hourly energy prices for the coming years. Two basic elements are required in order to create a PFC, (1) an estimate of the relative price structure, with an hourly resolution for the entire time period, and (2) current market values of forward products (EEX). It is important to emphasise that the PFC is not a forecast in the sense of an independent market evaluation, but rather a market-consistent breakdown of the market conditions to an hourly granularity. These two elements are supported by regression analysis, which represents the core elements in the creation of the PFC. The main input parameters are the calendar dates (weekdays, holidays, etc.) to




FIGURE 26.1 Price forward curve modelling approach.




FIGURE 26.2 A generated price forward curve (light grey) and the historically realised spot price (dark grey) compared.

obtain seasonality effects (intraday, weekly and yearly), historical meteorological data (daily temperature, wind and precipitation data) and the long-term market expectations. The PFC is then scaled so that the average price of the hourly PFC is equal to the future power contract at its corresponding time interval to obtain arbitrage-free conditions on weekly, monthly, quarterly and yearly peak and off-peak forward products. The calibration process is depicted in Figure 26.1. Figure 26.2 shows the hourly price Forward Curve created *a priori* over a monthly period together with the realised spot prices. The PFC is far smoother and less volatile than the realised spot prices since it is based on averaged historical data.

26.3.2 The Short-Term Process

The short-term process X_t is described by Equation (26.2). We have calibrated the volatility and the mean reversion to fit the distribution curve of the price difference between the Price Forward Curve and the realised spot price with EEX prices from January 1, 2011 until January 1, 2012. We have extracted the effect of price variations of the forward contracts by using daily updated PFCs to only include unexpected spot price variations due to short-term weather fluctuations and unexpected outages of power plants (PFC*). The empirical autocorrelation function η of the difference of the realized spot price and the PFC* can be approximated by the function η^τ , with τ the time lag in hours, where η was calibrated to

be 0.8. Then, with least-squares regression modelling, we estimate the mean reversion and the volatility parameters. Figure 26.3 depicts the comparison of the realised and the modelled price differences, i.e. (PFC*– realised spot price) and (PFC*–short-term spot price model).

We also investigated the variation of the short-term volatility as a function of peak and off-peak hours, week days and months (Figure 26.4). We found significant differences in the hourly and monthly structures, thus developing an empirical expression to model the short-term volatility as a function of the hour and the month. The parameters in the discrete jump process (26.3) were estimated by analysing the frequency and magnitude of the outliers. For this purpose a realised spot price outside the range of two standard deviations from the adjusted PFC* was defined as an outlier.




FIGURE 26.3 Comparison of the realised and modelled price differences: (PFC*–realised spot price) vs. (PFC*–short-term spot price model).




FIGURE 26.4 Realised short-term volatility as a function of hour and month.

26.3.3 The Long-Term Process

In order to reflect shifts in the level of spot prices due to changes in long-term forward contracts, we included daily shifts of the PFC based on a log-normal mean-reverting process of the front-year forward base contract. The parameters of the model are calibrated with 5 years of front-year forward contract data. The long-term process is expressed by Equation (26.4). Figure 26.5 shows a qq-plot of the realised front-year forward price and the forward model. The qq-plot suggests that forward model prices follow a distribution that fits the forward price cumulative distribution. Figure 26.6 depicts the realised front-year forward price and one scenario generated with the forward price model. The model seems to mimic the price behaviour of the realised prices. Figure 26.7 shows the corresponding comparison of the weekly log returns of realised forward prices and the path simulated by the forward price model.




FIGURE 26.5 qq-Plot of the realised front-year forward and forward model prices.




FIGURE 26.6 Realised front-year forward price and one generated scenario with the forward price model.




FIGURE 26.7 Weekly log returns of realised forward prices and one simulated path of the forward price model.




FIGURE 26.8 Term structure of volatility.

26.3.4 The Term Structure of Volatility

A set of generated spot price scenarios is used to evaluate the price of future options contracts available on the EEX. The distribution of the averaged spot price values for different delivery periods is used to evaluate the fair price of call and put options based on selected strike prices. These values are compared with the available quoted market prices. The long-term volatility in Equation (26.4) is calibrated so that calculated option prices for monthly, quarterly and annual contracts fall within the quoted bid–ask spread. A simplified scaling factor of the form $a \cdot t_d^{-\beta}$ is used to include the volatility term structure in the long-term volatility. The parameters are adjusted on a daily basis in order to be consistent with the market quotes from the brokers for the different delivery periods (see Figure 26.8). For comparison with

the implied volatility of yearly options, we note that the annualised volatility of the forward price model is of the order of 13–14%.

26.4 Scenario Generation and Valuation of Structured Products

The model defined above is used to generate multiple spot price paths. Figure 26.9 shows one path in comparison with the price forward curve, and Figure 26.10 shows a selection of five simulated paths. The model defined above is also used to generate multiple price paths for forward prices. [Figure 26.11](#) shows




FIGURE 26.9 Price forward curve (dark grey) and spot price model (light grey).




FIGURE 26.10 A sample set of spot price scenarios.




FIGURE 26.11 Cumulative forward price distribution.




FIGURE 26.12 Volatility surface for a yearly call option, Cal13 (quoted 20.04.2012).

the cumulative forward price distribution curve of 200 paths of a simulated front-year product for a day, week and monthly time period. Based on this information we can estimate the value at risk for different time periods.

26.4.1 Estimation of the Volatility Surface

Generated spot price scenarios are used to evaluate the price of option contracts available on the EEX. The distributions of the averaged spot price values for different delivery periods are used to evaluate the fair price of call and put options for different strike prices and times to expiry. Figure 26.12 shows the volatility surface of a call option on the yearly forward contract Cal13 based on quotes from 20.04.2012,




FIGURE 26.13 Estimated call option premiums for a yearly call option, Cal13 (quoted 20.04.2012).

when the underlying price was 51 Euro/MWh. Figure 26.13 shows that the calculated call option premiums are within the range of the bid–ask spread for an out-of-the-money option.

26.5 Structured Product Valuation

Non-standard financial contracts are products tailored to the needs of specific clients, but not sold in sufficient quantity for a standard price model to have been established. They are also referred to as exotic or structured products. A standard product usually trades actively and there is little uncertainty about its price so that the choice of model primarily affects how hedging is done. In the case of structured products, model risk is much greater because there is the potential for both pricing and hedging being impacted. Examples of structured products are swing options and virtual power plants (VPPs), see [Figure 26.14](#), which address the need of market participants to have flexibility in both the timing and volume of the energy delivered. The increasing variability of both generation (from solar and wind) and loads will also require more sophisticated and decentralised decision making. As a result of all these factors, interest in virtual power plants is gaining significant momentum within the energy industry. The price forward curve and spot price scenarios described in this paper are used to evaluate a fair price and value at risk for these non-standard products using the stochastic optimisation software TS-Energy (Zuur 2004, 2008, 2012).

26.5.1 Valuation and Risk Analysis: TS-Energy

TS-Energy is a valuation and risk analysis application for complex power assets and contracts. Its functionalities are as follows.

- Optimal use of storage capacity (e.g. hydro power, gas, wind capacity): When to sell, hold or buy (e.g. wind power at night); optimal operation, fair price, value at risk and bidding information.
- Optimal use and pricing of mid-term power contracts/VPPs.
- Pricing of contracts between suppliers of storage power and their customers.
- Arbitrage—taking advantage of today's forward pricing inefficiencies (brokers).




FIGURE 26.14 Simplified setup of a virtual power plant (VPP).

- Portfolio analysis and optimisation.
- Efficient planning and valuation of power generation plants and storage capacity. Expected net present value, value at risk and payback period of a business model.

The approach uses stochastic processes to describe uncertain inputs such as electricity prices and water influx to optimise assets without the requirement of perfect foresight imposed by standard linear programming software. TS-Energy uses a two-step process, a backwards integration step that determines an optimal operation strategy for an asset and a forwards integration step that applies this strategy to a set of scenarios to determine fair value and risk information for the asset (Zuur 2004, 2008, 2012, Helland and Winnington 2012). Hydropower assets are described by basins and operating units (generators and pumps) (see Figure 26.14), as well as by stochastic processes describing the net water influx for each basin. During the valuation step, the backward integration method is carried out in order to derive the action grid which holds the information about the best action at each settle date (e.g. to hold or exercise the option) given the market state (e.g. the underlying price) and the product state (e.g. the number exercises left of the swing contract or the volume of a basin). The modelling of the optimal dispatch strategy implies the determination of the value-maximising decisions for all states at each point in time in order to determine the right time for generation or pumping (see Figure 26.15). In the forwards integration step, this optimal dispatch strategy is applied to spot price scenarios generated by means of Monte-Carlo simulations. An adequate number of scenarios for convergence is in the range of 100 to 10,000, which can be used to derive the expected earnings, the probability distribution and the value at risk of the asset.

26.5.2 Virtual Power Plant Valuation

Based on the most recent PFC, and a calibrated set of weekly, hourly and jump processes, a set of price scenarios is generated using the spot price simulation model. These scenarios are used to evaluate the price of a VPP (virtual power plant) contract. These contracts are created and evaluated in TS-Energy.




FIGURE 26.15 Backward and forward process (TS-Energy 2010).

		Price in €/MWh																
		0.0	6.5	9.6	14.1	20.7	30.5	44.8	66.0	97.1	142.8	210.1	309.0	3000.0				
Delivery hour	1	-180	-180	-180	-180	-180	0	24	24	172	172	172	172	172				
	2	-180	-180	-180	-180	-180	-180	-156	24	24	204	204	204	204				
	3	-180	-180	-180	-180	-180	-180	-156	-156	24	24	204	204	204				
	4	-138	-138	-138	-138	-138	-138	-138	-114	-114	24	24	204	204				
	5	0	0	0	0	0	0	24	24	24	24	24	204	204				
	6	0	0	0	0	0	0	24	24	24	24	204	204	204				
	7	0	0	0	0	0	24	24	204	204	204	204	204	204				
	8	0	0	0	0	0	204	204	204	204	204	204	204	204				
	9	-180	-180	-180	-180	-180	24	24	204	204	204	204	204	204				
	10	-180	-180	-180	-180	-180	24	24	204	204	204	204	204	204				
	11	-180	-180	-180	-180	-180	24	204	204	204	204	204	204	204				
	12	-180	-180	-180	-180	0	24	204	204	204	204	204	204	204				
	13	-180	-180	-180	-180	-180	0	204	204	204	204	204	204	204				
	14	-180	-180	-180	-180	-180	0	24	204	204	204	204	204	204				
	15	-180	-180	-180	-180	-180	0	24	204	204	204	204	204	204				
	16	-180	-180	-180	-180	-180	-180	24	24	204	204	204	204	204				
	17	-180	-180	-180	-180	-180	-180	0	24	24	204	204	204	204				
	18	-180	-180	-180	-180	-180	-180	0	24	24	204	204	204	204				
	19	-180	-180	-180	-180	-180	0	24	24	24	204	204	204	204				
	20	-180	-180	-180	-180	-180	0	24	204	204	204	204	204	204				
	21	-180	-180	0	0	180	180	204	204	204	204	204	204	204				
	22	-180	-180	-180	0	0	100	100	157	157	157	157	157	157				
	23	-180	-180	-180	0	0	0	0	150	150	150	150	150	150				
	24	-180	-180	-180	-180	0	0	0	86	143	143	143	143	143				
		Volume in MWh																

FIGURE 26.16 Price-dependent hourly bids (TS-Energy 2010).

The distribution of the scenarios' end values is used to evaluate the fair price of the contract. Since there is no inflow to the VPP, the only stochastic variable is the spot price.

The calculations give the plant operator the hourly price boundaries for generation, holding and pumping in order to maximise net income until maturity by providing the optimal price-dependent hourly bids (see Figure 26.16). Figure 26.17 shows a typical result report of a plant simulation. It provides information about the fair price, the expected income from generation, the expected cost for pumping, the standard deviation and the Value at Risk of operating a hydropower plant for a given time period. Figure 26.18 depicts the cumulative distribution of the fair value of the VPP contract.

26.6 Conclusion

A stochastic spot price model has been developed in order to evaluate standard and off-standard products in the electrical market and for the optimisation of the power production in pump storage hydro parks. The spot price model is built up in several steps and calibration is carried out based on the price distribution curves of spot prices and the log returns of forward contracts, as well as standard option

	01.06.2012	31.08.2012	30.11.2012	01.03.2013	01.06.2013
Expected value	723'131	723'131	723'131	723'131	723'131
Expected price	723'131	669'815	497'454	286'380	0
Expected income (cumulated)	0	53'316	225'678	436'751	723'131
Expected income	0	53'316	172'362	211'073	286'380
Expected income by generation	0	250'967	812'522	1'534'952	2'465'884
Expected costs of pumping	0	197'651	586'845	1'098'201	1'742'753
Expected penalties	0	0	0	0	0
Expected costs of switching on/off	0	0	0	0	0
Standard deviation of value	0	31'253	64'642	102'054	129'977
Value at Risk (95%)	0	56'109	123'980	208'839	269'170
Generated electricity (MWh)	0	4'429	13'435	21'910	34'786
Consumed electricity (MWh)	0	6'667	18'821	30'030	46'381
Days generated	0	12	37	61	97
Days pumped	0	14	39	63	97
Expected volume (m ³)	0	571	680	612	0
Expected volume used for generation (m ³)	0	4'429	13'435	21'910	34'786
Expected volume used for pumping (m ³)	0	5'000	14'115	22'522	34'786

FIGURE 26.17 Valuation report (TS-Energy 2010).




FIGURE 26.18 Cumulative distribution of the VPP valuation.

contracts over several time periods. The price model is used to evaluate non-standard contracts, to optimise the scheduling of a hydro power pumped-storage plant, and for dynamical hedging of production.

Acknowledgements

The results of this paper were obtained in a joint development project of Axpo AG and Time-steps AG. The project was supported financially by the Axpo Trading Division. The authors wish to express their gratitude to Roman von Siebenthal, Ed Zuur, André Röthig and Jens Wartenburger for their valuable contributions.

References

- Helland, E. and Winnington, E., Stochastic optimization of the scheduling of a multi-stage pumped-storage hydro power plant. Report, Hydro scheduling in competitive markets, Bergen, Norway 2012.
- Müller, A., Schindmayr, G., Burger, M. and Klar, B., A spot market model for pricing derivatives in electricity markets. *Quant. Finance*, 2004, **4**, 109–122.
- Zuur, E., Ts-hydropower, approach document. Report, Time-steps AG, Switzerland, 2004.
- Zuur, E., Hydro 2008: calibration the mean reversion process. Report, Time-steps AG, Switzerland, 2008.
- Zuur, E., Ts-energy 2012, technical description. Report, Time-steps AG, Switzerland, 2012.

27

Modelling Spikes and Pricing Swing Options in Electricity Markets

Ben Hambly
Sam Howison
Tino Kluge

27.1	Introduction	585
27.2	Properties of the Model for Spot Prices	588
	The Spike Process • The Combined Process • Approximations	
27.3	Option Pricing.....	593
	Pricing Path-Independent Options • Pricing Options on Forwards • Pricing Options on Forwards With a Delivery Period	
27.4	Pricing Swing Options.....	600
	The Grid Approach • Numerical Results	
	References.....	605

Most electricity markets exhibit high volatilities and occasional distinctive price spikes, which result in demand for derivative products which protect the holder against high prices. In this paper we examine a simple spot price model that is the exponential of the sum of an Ornstein–Uhlenbeck and an independent mean-reverting pure jump process. We derive the moment generating function as well as various approximations to the probability density function of the logarithm of the spot price process at maturity T . Hence we are able to calibrate the model to the observed forward curve and present semi-analytic formulae for premia of path-independent options as well as approximations to call and put options on forward contracts with and without a delivery period. In order to price path-dependent options with multiple exercise rights like swing contracts a grid method is utilized which in turn uses approximations to the conditional density of the spot process.

Keywords: Energy derivatives; Financial mathematics; Stochastic jumps; Numerical methods for option pricing; Continuous time models; Derivative pricing models

27.1 Introduction

A distinctive feature of electricity markets is the formation of price spikes which are caused when the maximum supply and current demand are close, often when a generator or part of the distribution network fails unexpectedly. The occurrence of spikes has far reaching consequences for risk management and pricing purposes. In this context a parsimonious model with some degree of analytic tractability has clear advantages, and in this paper we propose and examine in detail a simple mean-reverting spot price process exhibiting spikes.

In our model the spot price process S is defined to be the exponential of the sum of three components: a deterministic periodic function f characterizing seasonality, an Ornstein–Uhlenbeck (OU) process X and a mean-reverting process with a jump component to incorporate spikes Y . We set this up formally. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space equipped with a filtration $(\mathcal{F}_t)$ satisfying the usual conditions. We let

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t Y_t &= -\beta Y_t dt + J_t dN_t, \end{aligned} \quad (27.1)$$

where N is a Poisson process with intensity λ and J is an independent identically distributed (i.i.d.) process representing the jump size, W is a Brownian motion and α , β , σ are constants. We assume W , N and J to be mutually independent processes adapted to the filtration.

Our model generalizes a number of earlier models. For example, the commonly used model of Lucia and Schwartz (2002) is obtained by setting $\beta = 0$ and taking $J = 0$. In this model, S_t is log-normally distributed giving analytic option price formulae very similar to those in the Black–Scholes model. To allow for a stochastic seasonality, a further component can be inserted into the model and, as long as this process has a normal distribution, analytic tractability is maintained. The main disadvantage of the Lucia and Schwartz models is their inability to mimic price spikes. To overcome this, jumps can be added to the model; for example, the case $\beta = 0$ in our model,

$$dX_t = -\alpha X_t dt + \sigma dW_t + J_t dN_t, \quad S_t = \exp(f(t) + X_t). \quad (27.2)$$

This model is briefly mentioned by Clewlow and Strickland (2000, Section 2.8). Analytic results are given by Deng (2000) based on transform analysis described by Duffie *et al.* (2000). Calibration to historical data and the observed forward curve is discussed by Cartea and Figueroa (2005) where practical results for the UK electricity market are given. More general versions are discussed by Benth *et al.* (2003) and applied to data from the Nordpool market. For these models to exhibit typical spikes the mean-reversion rate α must be extremely high, otherwise the jumps do not revert quickly enough.

Benth *et al.* (2007) introduce a set of independent pure mean-reverting jump processes of the form

$$S_t = \sum_i w_i Y_t^{(i)}, \quad dY_t^{(i)} = -\alpha_i Y_t^{(i)} dt + \sigma_i dL_t^{(i)},$$

where w_i are some positive weights and the $L^{(i)}$ are independent increasing càdlàg pure jump processes. The spot price process is a linear combination of the pure jump processes and, as there is no exponential function involved, positivity of the spot is achieved by allowing positive jumps only. An advantage of this formulation is that semi-analytic formulae for option prices on forwards with a delivery period can be derived. However, a full analysis of this class of models still seems to be in its early stages; there is some empirical work on fitting this model in the work of Klüppelberg *et al.* (2008).

There are a number of papers which discuss the stylized facts of electricity markets and seek to model them empirically. An empirical comparison of a number of models for the Californian market can be found in the work of Knittel and Roberts (2005). Other empirical approaches try to fit more advanced time series models (e.g. Koopman *et al.* 2007), or alternative mathematical models, such as jump diffusion and regime switching (e.g. Weron *et al.* 2004). It is clear that there is a need to capture seasonality and rapid price changes.

The model (27.1) we consider in this paper is an extension of (27.2), in which we allow for two different mean-reversion rates, one for the diffusive part and one for the jump part. The introduction of a mean-reverting spike process Y allows us to choose a higher mean-reversion rate β in order for the jump to

revert much more quickly and so mimic a price spike. This is useful for modelling the NordPool market where estimates of the mean-reversion speed are typically small, but this might not be needed in markets where the speed of mean reversion α is generally very high, like in the UKPX or EEX, where a jump diffusion with mean reversion may also be appropriate.

Returning to our model (27.1) and solving for X_t and Y_t , we have

$$\begin{aligned} X_t &= X_0 e^{-\alpha t} + \sigma \int_0^t e^{-\alpha(t-s)} dW_s, \\ Y_t &= Y_0 e^{-\beta t} + \sum_{i=1}^{N_t} e^{-\beta(t-\tau_i)} J_{\tau_i}, \end{aligned} \quad (27.3)$$

where τ_i indicates the random time of the occurrence of the i th jump. Thus, given $X_0 = x_0$, $X_t \sim N(x_0 e^{-\alpha t}, (\sigma^2/2\alpha)(1 - e^{-2\alpha t}))$. Properties of the spike process Y are not as obvious and will be examined in the following section. At this point we make no assumption on the jump size J but will later give results for exponentially and normally distributed jump sizes.

Note also that, although X and Y are both Markov processes, the price process S is not. We will therefore assume that all three components of the price process, i.e. the time t values $f(t)$, X_t and Y_t , are individually observable, implying we expect jumps not to be small.

Figure 27.1 shows a simulated sample path of the processes X , Y and the composed process S .

In Section 27.2 we derive important stochastic properties of the process, including the moment generating function and various approximations to its probability density function. Pricing of a variety of derivative contracts will be discussed in Sections 27.3 and 27.4, using the results obtained in Section 27.2.




FIGURE 27.1 Simulated sample paths of X , Y and S of the spot model (27.1). We use the following parameters which are of the same order of magnitude as calibrated values of the NordPool spot market, except the seasonality function $f(\cdot)$ which has been chosen arbitrarily: $f(t) = \ln(100) + 0.5\cos(2\pi t)$, $\alpha = 7$, $\sigma = 1.4$, $\beta = 200$, $J_t \sim \exp(1/\mu_J)$, $\mu_J = 0.4$, $\lambda = 4$.

27.2 Properties of the Model for Spot Prices

27.2.1 The Spike Process

The following result is known and given by Duffie *et al.* (2000) in a more general framework.

Lemma 2.1: Moment Generating Function of the Spike Process, Y_t

Let $\{J_1, J_2, \dots\}$ be a series of i.i.d. random variables. We assume that there is a $\theta_0 > 0$ such that the moment generating function $\Phi_J(\theta) := \mathbb{E}[e^{\theta J}]$ exists for $\theta < \theta_0$. Let $\{\tau_1, \tau_2, \dots\}$ be the random jump times of a Poisson process N with intensity λ . Then the process Y with initial condition $Y_0 = 0$ has, for all $\theta < \theta_0$, $t \geq 0$, the moment generating function

$$\Phi_Y(\theta, t) := \mathbb{E}[e^{\theta Y_t}] = \exp\left(\lambda \int_0^t \Phi_J(\theta e^{-\beta s}) - 1 ds\right). \quad (27.4)$$

Furthermore, the first two moments of Y_t are given by

$$\begin{aligned} \mathbb{E}[Y_t] &= \Phi_Y'(0, t) = \frac{\lambda}{\beta} \mathbb{E}[J](1 - e^{-\beta t}), \\ \mathbb{E}[Y_t^2] &= \Phi_Y''(0, t) = \mathbb{E}[Y_t]^2 + \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}), \end{aligned}$$

and in particular we have

$$\text{var}[Y_t] = \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}).$$

Remark 1 (asymptotics for $\beta \rightarrow \infty$): As remarked above, in practice the timescale $1/\beta$ for mean reversion of spikes is much shorter than any of: the contract lifetime T ; the diffusive mean-reversion time $1/\alpha$; the volatility timescale $1/\sigma^2$; and the mean arrival time of spikes $1/\lambda$. We therefore calculate approximations for the moment generating function of the spike process, and for its distribution, as $\beta \rightarrow \infty$.

To analyse the behaviour of the moment generating function for large β we make the substitution $u = \theta e^{-\beta s}$ in the integrand to obtain

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \int_{\theta e^{-\beta t}}^{\theta} \frac{\Phi_J(u) - 1}{u} du\right).$$

For fixed θ , t , as $\beta \rightarrow \infty$ we have

$$\int_0^{\theta e^{-\beta t}} \frac{\Phi_J(u) - 1}{u} du = \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t}),$$

because $\Phi_J(u) = 1 + \mathbb{E}[J]u + O(u^2)$, $u \rightarrow 0$, and so

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \left(\int_0^{\theta} \frac{\Phi_J(u) - 1}{u} du - \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t}) \right)\right). \quad (27.5)$$

Example 2.2: Exponentially Distributed Jump Size

If $J \sim \text{Exp}(1/\mu_J)$ with mean jump size μ_J , then $\Phi_J(\theta) = 1/(1 - \mu_J)$ exists for $\theta < \theta_0 = 1/\mu_J$. We obtain

$$\Phi_Y(\theta, t) = \left(\frac{1 - \theta \mu_J e^{-\beta t}}{1 - \theta \mu_J} \right)^{\lambda/\beta}, \quad \theta < 1/\mu_J$$

As $t \rightarrow \infty$, we have $\Phi_Y(\theta, t) \rightarrow (1 - \theta \mu_J)^{-\lambda/\beta}$ for $\theta < 1/\mu_J$ so the stationary distribution for Y is the Gamma distribution $\text{Gamma}(\lambda/\beta, 1/\mu_J)$. As $\beta \rightarrow \infty$, we also have $\Phi_Y(\theta, t) = 1 + \theta \mu_J \lambda/\beta + O(\beta^{-2}) \approx (1 - \theta \mu_J)^{-\lambda/\beta}$ for $\theta < 1/\mu_J$. Thus Y_t is distributed approximately as $\text{Gamma}(\lambda/\beta, 1/\mu_J)$ for large β .

The mean and variance of the spike process Y_t with $Y_0 = 0$ are

$$\mathbb{E}[Y_t] = \frac{\lambda \mu_J}{\beta} (1 - e^{-\beta t}), \quad \text{var}[Y_t] = \frac{\lambda \mu_J^2}{\beta} (1 - e^{-2\beta t}).$$

27.2.2 The Combined Process

Having examined the properties of the spike process Y we conclude properties of the sum $X_t + Y_t$ and consequently of the price $S_t = \exp(f(t) + X_t + Y_t)$.

Theorem 2.3:

Let the spot process S be defined by (27.1) and let $Z_t := \ln S_t = f(t) + X_t + Y_t$ with X_0 and Y_0 given. The moment generating function of Z_t exists for $\theta < \theta_0$ and is given by

$$\begin{aligned} \mathbb{E}e^{\theta Z_t} = \exp & \left(\theta f(t) + \theta X_0 e^{-\alpha t} + \theta^2 \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha t}) \right. \\ & \left. + \theta Y_0 e^{-\beta t} + \lambda \int_0^t \Phi_J(\theta e^{-\beta s}) - 1 ds \right) \end{aligned} \quad (27.6)$$

Proof: The processes X and Y are independent so the expectation of the product is the product of the expectations. The moment generating function of Y_t is given in Lemma 2.1 which yields the result. ■

If $\theta_0 > 1$, the expectation value of the spot process at time t , S_t , immediately follows by setting $\theta = 1$. We note that, in general, the moments for the spot price $\mathbb{E}S_t^\theta$ only exist for $\theta < \theta_0$. For instance, if the jump distribution is exponential with mean μ_J , the spot price has less than $1/\mu_J$ moments.

27.2.3 Approximations

We now derive approximations to the density functions of the spike process at maturity T for large mean-reversion values β of the spike process. Although we can always compute the density by Laplace inversion of the moment generating function, an explicit expression for the density allows for more efficient algorithms and more explicit option pricing formulae. Here we only provide an expression for the density of a ‘truncated’ spike process $\tilde{Y}_T$ as defined below. However, knowledge of the density of $\tilde{Y}_T$ alone will help us to efficiently construct a grid to price swing options. We start by defining the truncated spike process $\tilde{Y}$, showing that $\tilde{Y}_t$ provides a good approximation to the value $\tilde{Y}_t$ for large values of β , deriving a general formula for the density function of $\tilde{Y}_t$ and finally making it more explicit by considering an exponential jump size distribution.

For very high mean-reversion rates β and small jump intensities λ , the dominant contribution to the density of the spike process comes from the last jump. We therefore introduce the truncated spike process

$$\tilde{Y}_t := \begin{cases} J_{N_t} e^{-\beta(t-\tau_{N_t})}, & N_t > 0, \\ 0, & N_t = 0. \end{cases} \quad (27.7)$$

Note that we only consider Y starting from 0; any other starting point can be incorporated by adding the initial value.

Lemma 2.4:

$\tilde{Y}_t$ is identically distributed as

$$Z_t := \begin{cases} J_1 e^{-\beta\tau_1}, & \tau_1 \leq t, \\ 0, & \tau_1 > t. \end{cases}$$

Proof: We use the reversibility property that if $N = \{N_t; t \in \mathbb{R}_+\}$ is a Poisson process, then $\hat{N} = \{-N_{-t}; t \in \mathbb{R}_+\}$ is also a Poisson process. As τ_{N_t} is the jump time of the last jump before t , this translates into the first jump of the reversed process and hence $t - \tau_{N_t}$ and τ_1 are identically distributed, given $N_t > 0$. If $N_t = 0$, then there has been no jump in $[0, t]$ and the same applies for the reversed process and so this is equivalent to $\tau_1 > t$. ■

Lemma 2.5 (moment generating function of the truncated spike process):

The random variable $\tilde{Y}_t$ of the truncated spike process at time t with initial condition $\tilde{Y}_0 = 0$ has a moment generating function for $\theta < \theta_0$ which is given by

$$\Phi_{\tilde{Y}_t}(\theta) = 1 + \lambda \int_0^t (\Phi_J(\theta e^{-\beta s}) - 1) e^{-\lambda s} ds.$$

The first two moments are given by

$$\begin{aligned} \mathbb{E}[\tilde{Y}_t] &= \frac{\lambda}{\beta + \lambda} \mathbb{E}[J](1 - e^{-(\beta + \lambda)t}), \\ \mathbb{E}[\tilde{Y}_t^2] &= \frac{\lambda}{2\beta + \lambda} \mathbb{E}[J^2](1 - e^{-(2\beta + \lambda)t}). \end{aligned}$$

Proof: By Lemma 2.4 we only need to determine the moment generating function of

$$Z_t := J e^{-\beta\tau} I_{\tau \leq t}, \quad \tau \sim \text{Exp}(\lambda),$$

where I_A denotes the indicator function for the event A . Given the jump time τ we have

$$\mathbb{E}[e^{\theta Z_t} \mid \tau = s] = \Phi_J(\theta e^{-\beta s} I_{s \leq t}),$$

and so

$$\begin{aligned}\mathbb{E}[e^{\theta Z_t}] &= \mathbb{E}[\mathbb{E}[e^{\theta Z_t} | \tau]] \\ &= \int_0^\infty \Phi_J(\theta e^{-\beta s} I_{s \leq t}) \lambda e^{-\lambda s} ds \\ &= \int_0^t \Phi_J(\theta e^{-\beta s}) \lambda e^{-\lambda s} ds + e^{-\lambda t}.\end{aligned}$$

The first two moments are given by $\mathbb{E}[\tilde{Y}_t] = \Phi'_{Y_t}(0)$ and $\mathbb{E}[\tilde{Y}_t^2] = \Phi''_{Y_t}(0)$. ■

Remark 2 (pointwise convergence of the moment generating functions): The moment generating function of the truncated spike process converges pointwise to the moment generating function of the spike process for either $\lambda \rightarrow 0$ or $\beta \rightarrow \infty$ with t and θ fixed. First consider $\lambda \rightarrow 0$. Fix all other parameters and set $g(s; \beta, \theta) := \Phi_J(\theta e^{-\beta s}) - 1$, then

$$\begin{aligned}\Phi_{Y_t}(\theta) &= \exp\left(\lambda \int_0^t g(s; \beta, \theta) ds\right) \\ &= 1 + \lambda \int_0^t g(s; \beta, \theta) ds + O(\lambda^2), \\ \Phi_{\tilde{Y}_t}(\theta) &= 1 + \lambda \int_0^t g(s; \beta, \theta) e^{-\lambda s} ds \\ &= 1 + \lambda \int_0^t g(s; \beta, \theta) ds + O(\lambda^2).\end{aligned}$$

To see the convergence for $\beta \rightarrow \infty$ with λ , t and θ fixed, first note that from (27.5) we have

$$\Phi_{Y_t}(\theta) = 1 + \frac{\lambda}{\beta} \int_0^\theta \frac{\Phi_J(u) - 1}{u} du + O(1/\beta^2).$$

Also from lemma 2.5,

$$\begin{aligned}\Phi_{\tilde{Y}_t}(\theta) &= 1 + \lambda \int_0^t (\Phi_J(\theta e^{-\beta s}) - 1) e^{-\lambda s} ds \\ &= 1 + \frac{\lambda}{\beta} \int_{\theta e^{-\beta t}}^\theta \frac{\Phi_J(u) - 1}{u} \left(\frac{u}{\theta}\right)^{\lambda/\beta} du\end{aligned}$$

by setting $\theta e^{-\beta s} = u$. Now as $\beta \rightarrow \infty$, $(u/\theta)^{\lambda/\beta} \rightarrow 1$ except in a small region $u = O(\theta e^{-\beta/\lambda})$, which makes a negligible (exponentially small) contribution to the integral. Likewise, we may replace the lower limit of integration by 0 and incur a similarly small error. Hence,

$$\begin{aligned}\Phi_{\tilde{Y}_t}(\theta) &= 1 + \frac{\lambda}{\beta} \int_0^\theta \frac{\Phi_J(u) - 1}{u} du + o\left(\frac{1}{\beta}\right) \\ &= \Phi_{Y_t}(\theta) + o\left(\frac{1}{\beta}\right).\end{aligned}$$

Two examples of the approximated and exact moment generating function using our standard parameters can be seen in [Figure 27.2](#).

Lemma 2.6 (distribution of the truncated spike process):

Let the jump size distribution have density function f_J . Then the truncated spike process $\tilde{Y}_t$ as defined above has the cumulative distribution function




FIGURE 27.2 Moment generating function of Y_t and $\tilde{Y}_t$ denoted by mgf and approximation, respectively. On the left we use $J \sim \text{Exp}(1/\mu_J)$ and on the right $J \sim N(\mu_J, \mu_J^2)$. The same parameters as in Figure 27.1 are used and we set $t = 1$.

$$F_{\tilde{Y}_t}(x) = e^{-\lambda t} I_{x \geq 0} + \int_{-\infty}^x f_{\tilde{Y}_t}(y) dy \quad t \geq 0,$$

with probability density function

$$f_{\tilde{Y}_t}(x) = \frac{\lambda}{\beta} \frac{1}{|x|^{1-\lambda/\beta}} \left| \int_x^{xe^{\beta t}} f_{J_t}(y) |y|^{-\lambda/\beta} dy \right|, \quad x \neq 0. \quad (27.8)$$

Proof: Based on Lemma 2.4 it suffices to determine the distribution of

$$\tilde{Y}_t = JZI_{\tau \leq t}, \quad Z := e^{-\beta \tau}, \quad \tau \sim \text{Exp}(\lambda).$$

It follows that Z is the (β/λ) th power of a uniformly distributed random variable on $[0, 1]$ and its density is given by

$$f_z(x) = \frac{\lambda}{\beta} x^{-(1-(\lambda/\beta))} I_{x \in [0, 1]}.$$

As $\mathbb{P}(\tau > t) = e^{-\lambda t}$ we obtain the cdf of $ZI_{\tau \leq t}$ as

$$\begin{aligned} F_{ZI_{\tau \leq t}}(x) &= e^{-\lambda t} I_{x \geq 0} + \int_{-\infty}^x f_{ZI_{\tau \leq t}}(y) dy, \\ f_{ZI_{\tau \leq t}}(x) &= \frac{\lambda}{\beta} x^{-(1-(\lambda/\beta))} I_{x \in [e^{-\beta t}, 1]}, \end{aligned}$$

and the distribution of the product of two independent random variables J and $ZI_{\tau \leq t}$ is then given by

$$\begin{aligned} F_{JZI_{\tau \leq t}}(c) &= e^{-\lambda t} I_{c \geq 0} + \int_{-\infty}^c f_{JZI_{\tau \leq t}}(x) dx, \\ f_{JZI_{\tau \leq t}}(c) &= \int_{-\infty}^{\infty} f_{ZI_{\tau \leq t}}(c/x) \frac{f_J(x)}{|x|} dx. \end{aligned}$$

With

$$\begin{aligned} f_{JZI_{\tau \leq t}}(c/x) &= \frac{\lambda}{\beta} \frac{1}{c^{1-(\lambda/\beta)}} I_{x \in [c, ce^{\beta t}]} x^{1-(\lambda/\beta)}, \quad c > 0, \\ f_{JZI_{\tau \leq t}}(c/x) &= \frac{\lambda}{\beta} \frac{1}{|c|^{1-(\lambda/\beta)}} I_{x \in [ce^{\beta t}, c]} |x|^{1-(\lambda/\beta)}, \quad c < 0, \end{aligned}$$

the desired result follows. ■




FIGURE 27.3 Distribution of the spike process (Y_t) at T with a jump size of $J \sim \text{Exp}(1/\mu_J)$. We use approximation (27.9) and compare it with the exact density as produced by a Monte-Carlo simulation. We use the same parameters as in Figure 27.1.

Example 2.7 (exponential jump size distribution):

Let $J \sim \text{Exp}(1/\mu_J)$ be exponentially distributed. The probability density function for the truncated spike process $\tilde{Y}_t$ is

$$f_{Y_t}(x) = \frac{\lambda}{\beta \mu_J^{\lambda/\beta}} \frac{\Gamma(1 - \lambda/\beta, x/\mu_J) - \Gamma(1 - \lambda/\beta, x e^{\beta t}/\mu_J)}{x^{1-\lambda/\beta}}, \quad x > 0, \quad (27.9)$$

where $\Gamma(a, x)$ is the incomplete Gamma function. The approximation is a good fit to the exact density for typical market parameters as can be seen in Figure 27.3. The only discrepancy occurs at $Y_t = 0$ where the density has a singularity. We use this approximation in Section 27.4.1 to efficiently generate a grid to price swing options.

27.3 Option Pricing

The electricity market with the model presented is obviously incomplete. Not only are we faced with a non-hedgeable jump risk but also we cannot use the underlying process (S_t) to hedge derivatives due to inefficiencies in storing electricity. From now on we assume the model is specified in the risk-neutral measure $\mathbb{Q}$ as

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_t dt + J_t dN_t, \end{aligned} \quad (27.10)$$

where W is a Brownian motion under $\mathbb{Q}$ and N a Poisson process with intensity λ under $\mathbb{Q}$. For simplicity of notation we use the same parameters as in (27.1) but note that they might differ from the parameters under the real-world measure. A common modelling assumption is that the risk-neutral model has a jump structure that is similar to that observed under $\mathbb{P}$. This is certainly true for a discrete jump size distribution, so we will assume that it is the case here.

Lemma 3.1 (seasonal function consistent with the forward curve):

Let $t = 0$ and $F_0^{[T]}$ be the forward at time 0 maturing at time T ; then the risk-neutral seasonality function is given by

$$f(T) = \ln F_0^{[T]} - X_0 e^{-\alpha T} - Y_0 e^{-\beta T} - \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha T}) - \lambda \int_0^T \Phi_j(e^{-\beta s}) - 1 \, ds. \quad (27.11)$$

Proof: The forward price is $F_0^{[T]} = \mathbb{E}^Q[S_T]$ and so the result follows from (27.6). ■

This result forms an important part of the calibration of the model. As the model is incomplete the calibration procedure depends on the set of liquid derivatives used. Here we assume a continuous forward curve is observable in the market, i.e. values of $F_0^{[T]}$ are given. This is not a realistic assumption but there are ways to generate a continuous curve consistent with discretely observed prices; see Figure 27.4 and Kluge (2006, Section 2.2.2.5) for more details.

For the sake of simplicity we adopt a policy of choosing parameters from the real-world measure $\mathbb{P}$ if they are not uniquely determined by the set of observed derivative prices. This is equivalent to saying we choose a risk-neutral measure $\mathbb{Q}$ which changes as few parameters of the model as possible. As there are at least five parameters ($\alpha, \beta, \sigma, \lambda, \mu_j$) and seasonality to fit we need to trade-off those calculated using historical data and the observed derivative prices. The volatility parameter σ remains unchanged by any equivalent measure change so we can always determine it from historical data. If we only see forward prices in the market we can also calibrate all other parameters to historical data except the seasonal function. In brief, this can be done by de-seasonalizing the data using the real-world seasonality and then making a first estimation of α from which we can start filtering suspected spikes. From the reduced dataset, α can be re-estimated and suspected spikes filtered recursively. Having determined all parameters from historical data we finally calculate the risk-neutral seasonality function f from the observed forward curve based on (27.11).




FIGURE 27.4 Interpolation of the forward curve by a seasonal function and spline correction. Three years worth of spot history data has been used to calibrate a seasonality function which is then used as a first approximation of the forward curve. The difference between the seasonality function and the observed forward prices is then corrected by a piecewise quadratic polynomial.

27.3.1 Pricing Path-Independent Options

If the pay-off of an option on the spot at maturity T is given by $g(S_T)$ then its arbitrage free price at time t is given by

$$V(x, y, t) = e^{-r(T-t)} \mathbb{E}^Q[g(S_T) | X_t = x, Y_t = y].$$

Although we do not have an expression for the density of S_T we know its moment generating function and so can apply Laplace transform methods to calculate the expectation value. For an overview, see Cont and Tankov (2004, Section 11.1.3)* or Carr and Madan (1998) and Lewis (2001). Consider, for example, put or call options. Let $Z_t = \ln S_t$ and let $\Phi_{Z_t}(\theta)$ be its moment generating function, as given in (27.6). Now define its truncated moment generating function by

$$G_v(x, t) := \mathbb{E}[e^{vZ_t} I_{\{Z_t \leq x\}}] = \int_{-\infty}^x e^{vy} dF_{Z_t}(y),$$

which can be computed using a generalization of Lévy's inversion theorem:

$$G_v(x) = \frac{\Phi_{Z_t}(v)}{2} - \frac{1}{\pi} \int_0^\infty \frac{\Im(\Phi_{Z_t}(v + i\theta)e^{-i\theta x})}{\theta} d\theta.$$

The price of a put option is then

$$\begin{aligned} \mathbb{E}[(K - S_T)^+] &= K \mathbb{E}[I_{S_T \leq K}] - \mathbb{E}[S_T I_{S_T \leq K}] \\ &= KG_0(\ln K) - G_1(\ln K), \end{aligned}$$

and by put–call parity we obtain the price of a call option.

27.3.2 Pricing Options on Forwards

For a forward contract at time t , understood to be today, maturing at T the strike of a zero-cost forward is given by

$$F_t^{[T]} = \mathbb{E}^Q[S_T | \mathcal{F}_t].$$

The most common options on forwards are puts or calls maturing at the same time as the underlying forward, i.e. the pay-off is given by $(F_t^{[T]} - K)^+$, which is equivalent to $(S_T - K)^+$. We can price these contracts based on the dynamics of the spot and using methods developed above. However, by analysing the dynamics of the forward curve implied by the spot price model we will gain further insights and be able to relate the price of an option to the Black-76 formula (Black 1976), which is still widely used in practice.

Recall that the expectation value of S_T is equal to the moment generating function given in (27.6) at $\theta = 1$. For $F_t^{[T]} = \mathbb{E}^Q[S_T | X_t, Y_t]$ we obtain

$$F_t^{[T]} = \exp \left(f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_f(e^{-\beta s}) - 1 ds \right). \quad (27.12)$$

* They describe the method in terms of a complex valued characteristic function and Fourier inversion, but by allowing complex values the method can also be written in terms of Laplace transforms.

For fixed T , the dynamics of the forward maturing at T is then

$$\begin{aligned} \frac{dF_t^{[T]}}{F_t^{[T]}} = & -\lambda(\Phi_J(e^{-\beta(T-t)})-1)dt + \sigma e^{-\alpha(T-t)} dW_t \\ & + (\exp(J_t e^{-\beta(T-t)})-1) dN_t. \end{aligned} \quad (27.13)$$

The forward is a martingale under Q by definition, and so the drift term compensates the jump process. For large time to maturities $T-t$, a jump in the underlying process has only very limited effect on the forward. More precisely, if the relative change in the underlying is $\exp(J_t)-1$ the forward changes relatively by $\exp(J_t e^{-\beta(T-t)})-1$. In addition to the jump component the dynamics follows a deterministic volatility process starting with a low volatility $\sigma e^{-\alpha T}$ at $t=0$ and increasing to σ at maturity. Without the jump component there are clear similarities with the Black-76 model.

For pricing purposes we need to find the distribution of $F_T^{[T]}$ in terms of its initial condition $F_t^{[T]}$. We have

$$\begin{aligned} \ln F_T^{[T]} &= f(T) + X_T + Y_T, \\ \ln F_t^{[T]} &= f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) \\ &\quad + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds. \end{aligned} \quad (27.14)$$

Eliminating the seasonality component $f(T)$, and using the relations

$$\begin{aligned} X_T - X_t e^{-\alpha(T-t)} &= \sigma \int_t^T e^{-\alpha(T-s)} dW_s, \\ Y_T - Y_t e^{-\beta(T-t)} &= \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)}, \end{aligned}$$

we finally get

$$\begin{aligned} \ln F_T^{[T]} - \ln F_t^{[T]} &= \sigma \int_t^T e^{-\alpha(T-s)} dW_s + \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)} \\ &\quad + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds. \end{aligned}$$

Without the jump component, $F_T^{[T]}$ would be log-normally distributed. In order to relate the pricing of options to the Black-76 formula even in the presence of spike risks, we assume that $F_T^{[T]}$ is log-normally distributed in a first approximation. We ignore the heavy tails caused by the spike risk and so expect to underestimate prices of far out-of-the-money calls but should do well with at-the-money calls.

We define the approximation by matching the first two moments but take into account that by definition $F_t^{[T]}$ is a martingale for a fixed maturity T and in order to keep the same property we set

$$\ln F_T^{[T]} \approx \ln F_t^{[T]} + \xi, \quad \xi \sim N\left(-\frac{1}{2}\hat{\sigma}^2(T-t), \hat{\sigma}^2(T-t)\right),$$

and $\hat{\sigma}^2(T-t) := \text{var}[\ln F_T^{[T]} | F_t]$, i.e.

$$\begin{aligned} \hat{\sigma}^2(T-t) &= \text{var}\left[\sigma \int_t^T e^{-\alpha(T-s)} dW_s + \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)}\right] \\ &= \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha(T-t)}) + \frac{\lambda}{2\beta} \mathbb{E}[J^2] (1 - e^{-2\beta(T-t)}). \end{aligned}$$

Remark 1 (term structure of implied volatility):

Comparing this result with the setting of Black-76 (Black 1976) where $dF = F\sigma dW$ and so $F_T = F_t \exp(\xi)$ with $\xi \sim N(-\frac{1}{2}\sigma^2(T-t), \sigma^2(T-t))$, we conclude that $\hat{\sigma}$ is the implied Black-76 volatility and in a first approximation given by

$$\hat{\sigma}^2 \approx \frac{(\sigma^2 / 2\alpha)(1 - e^{-2\alpha(T-t)}) + (\lambda / 2\beta)\mathbb{E}[J^2](1 - e^{-2\beta(T-t)})}{T-t}, \quad (27.15)$$

which is shown in Figure 27.5. It can be seen that the spike process has a much more significant impact on the implied volatility for short maturities rather than for long-term maturities. As far as the price of an at-the-money call is concerned, the additional jump risk adds an almost constant premium to the price to be paid without any jump risk.

Remark 2 (implied volatility across strikes):

The approximation does not predict a change of implied volatility across strikes. However, the jump risk introduces a skew as can be seen in Figure 27.6 where the exact solution based on Section 27.3.1 has been used to calculate implied volatilities. The bigger the mean jump size and hence the bigger $\mathbb{E}[J_2]$, the more profound is the skew.

27.3.3 Pricing Options on Forwards With a Delivery Period

As electricity is a flow variable, forwards always specify a delivery period. The results of the previous section can therefore only be seen as an approximation to option prices on forwards with short delivery periods, like one day. Here we only consider options on forwards maturing at the beginning of the delivery period, i.e. the pay-off is given by some function of $F_t^{[T_1, T_2]}$ at time T_1 . An option on such a forward is conceptually similar to an Asian option in the Black–Scholes world. One method of pricing Asian options is to approximate the distribution of the integral by a log-normal distribution and this can be done by matching the first two moments; see Turnbull and Wakeman (1991), for example. Once the parameters of the approximate log-normal distribution have been determined, pricing options comes down to pricing in the Black–Scholes or Black-76 setting.

The strike price of a zero cost forward with a delivery period is generally given by a weighted average of instantaneous forwards of the form

$$F_t^{[T_1, T_2]} = \int_{T_1}^{T_2} w(T; T_1, T_2) F_t^{[T]} dT,$$




FIGURE 27.5 Implied volatilities and prices. The left graph shows implied volatilities with respect to time to maturity where approximation (27.15) is used. The three lines correspond to no jumps ($\mu_j = 0$), small jumps ($\mu_j = 0.4$) and big jumps ($\mu_j = 0.8$). In the right graph the corresponding prices of an at-the-money call are plotted. Parameters are $r = \ln(1.05)$, $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4.0$, $F_0^{[T]} = 100$, $K = 100$. Note, for an exponential distributed jump size J we have $\mathbb{E}[J^2] = 2\mu_j^2$.




FIGURE 27.6 Implied volatilities across strikes and sample paths. The upper graph shows the implied volatility for one maturity $T = 0.2$ based on the exact solution. The approximate solution (27.15) yields 0.82 and 0.85 for the small and big jumps, respectively. Sample paths of the model with the same parameters are drawn in the lower two graphs, where the left path is generated with a low mean jump size ($\mu_j = 0.4$) and the right with a high mean jump size ($\mu_j = 0.8$). All the other parameters are the same as in [Figure 27.5](#).

where for a settlement at maturity T_2 the weighting factor is given by $w(T; T_1, T_2) = 1/(T_2 - T_1)$ and for instantaneous settlement the discounting alters the weighting to $w(T; T_1, T_2) = re^{-rT} / (e^{-rT_1} - e^{-rT_2})$.

The second moment of $F_{T_1}^{[T_1, T_2]}$ is given by

$$\begin{aligned} & \mathbb{E}^Q \left[\left(\int_{T_1}^{T_2} w(T) F_{T_1}^{[T]} dT \right)^2 \middle| \mathcal{F}_t \right] \\ &= \int_{T_1}^{T_2} \int_{T_1}^{T_2} w(T) w(T^*) \mathbb{E}^Q \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right] dT dT^*, \end{aligned}$$

and the expectation of the product of two individual forwards $\mathbb{E}^Q[F_{T_1}^{[T]} F_{T_1}^{[T^*]} | \mathcal{F}_t]$ can be derived using the solution of the forward (27.12) as follows:

$$\begin{aligned} \ln F_{T_1}^{[T]} &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{\beta(T-T_1)} \\ &\quad \times \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad + \lambda \int_0^{T-T_1} \Phi_J(e^{-\beta s}) - 1 ds - \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds \\ &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{\beta(T-T_1)} \\ &\quad \times \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad - \lambda \int_0^{T_1-t} \Phi_J(e^{-\beta(T-T_1)} e^{-\beta s}) - 1 ds, \end{aligned}$$

and so

$$\begin{aligned}
 \ln F_t^{[T]} &= \ln F_t^{[T^*]} \\
 &= \ln F_t^{[T]} + \ln F_t^{[T^*]} (e^{-\alpha(T-T_1)} + e^{-\alpha(T^*-T_1)}) \sigma \\
 &\quad \times \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + (e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)}) \\
 &\quad \times \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} - \frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)}) \\
 &\quad \times (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) - \ln \Phi_Y(e^{-\beta(T-T_1)}) \\
 &\quad - \ln \Phi_Y(e^{-\beta(T^*-T_1)}),
 \end{aligned}$$

which gives

$$\begin{aligned}
 \mathbb{E}^Q[F_t^{[T]} F_t^{[T^*]} | \mathcal{F}_t] &= \mathbb{E}^Q[\exp(\ln F_t^{[T]} + F_t^{[T^*]} | \mathcal{F}_t)] \\
 &= F_t^{[T]} F_t^{[T^*]} \frac{\Phi_Y(e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)})}{\Phi_Y(e^{-\beta(T-T_1)}) \Phi_Y(e^{-\beta(T^*-T_1)})} \\
 &\quad \times \exp\left(-\frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}))\right) \\
 &\quad \times \exp\left(\frac{\sigma^2}{4\alpha} (1 + e^{-\alpha(T^*-T)})^2 (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)})\right).
 \end{aligned}$$

How well the moment matching procedure works is shown in Figure 27.7 where the density of a forward $F_{T_1}^{[T_1, T_2]}$ is compared with the density obtain by the approximation. The shapes of both densities are similar but differences in values are clearly visible. As a result, one might not expect call option prices based on the approximate distribution to be very close to the exact prices for all strikes K but still close enough to be useful. For an at-the-money strike call option, prices for varying maturities are shown in Figures 27.8 and 27.9. As it turns out, the approximation gives very good results for short delivery periods and is still within a 5% range for delivery periods of one year.




FIGURE 27.7 Distribution of $F_{T_1}^{[T_1, T_2]}$. Parameters are as before, see Figure 27.5, and $\mu_j = 0.4$, $T_1 = 1$, $T_2 = 1.25$. The red curve is based on a Monte-Carlo simulation with 10^6 sample paths; the green curve is the log-normal approximation matching the first two moments.




FIGURE 27.8 Value of an at-the-money call option on a forward $F_{T_1}^{[T_1, T_2]}$ depending on the delivery period $T_2 - T_1$. Parameters are as before, see Figure 27.5, and $\mu_j = 0.4$, $T_1 = 1$, $K = 100$. The Monte-Carlo result is based on 100,000 sample paths for each duration. Note, any forward $F_{T_1}^{[T_1, T_2]}$ delivers exactly 1 MWh over the delivery period $[T_1, T_2]$.




FIGURE 27.9 Same as in Figure 27.8 but where the volume of the forward is proportional to the delivery period, i.e. we assume a constant consumption of 1 MWh per year, which is about 114 W. Here, the price is simply $T_2 - T_1$ times the price of a standard call option on $F_{T_1}^{[T_1, T_2]}$.

27.4 Pricing Swing Options

Swing contracts are a broad class of path-dependent options allowing the holder to exercise a certain right multiple times over a specified period but only one right at a time* or per time interval like a day. Such a right could be the right to receive the pay-off of a call option. Other possibilities include a mixture of different pay-off functions like calls and puts or calls with different strikes. Another very common feature is to allow the holder to exercise a multiple of a call or put option at once, where the multiple is called volume. This generally involves further restriction on the volume, like upper and lower bounds for each right and for the sum of all trades.

Swing contracts can be seen as an insurance for the holder against excessive rises in electricity prices. Assuming the prices generally revert to a long-term mean, even a small number N of exercise opportunities suffices to cover the main risks and hence make the premium of the contract cheaper. Sometimes,

* This will also involve a 'refraction period' in which no further right can be exercised.

swing contracts are bundled with forward contracts. The forward contract then supplies the holder with a constant stream of energy at a fixed pre-determined price. If the strike price of the call options of the swing contract is set to the forward price, the swing contract will allow for flexibility in the volume the customer receives for the fixed price. They can either 'swing up' or 'swing down' the volume of energy and hence the name swing contract. One cannot assume that the holder always exercises the contract in an optimal way to maximize expected profit but they might also exercise according to their own internal energy demands.

It is only very recently that articles on numerical pricing methods for swing options have appeared in the literature. We can identify a few main approaches, all based on dynamic programming principles. A Monte-Carlo method and ideas of duality theory are utilized by Meinshausen and Hambly (2004) to derive lower and upper bounds for swing option prices. The main advantages of the method being its flexibility, as it can be easily adapted to any stochastic model of the underlying, and its ability to produce confidence intervals of the price. Monte-Carlo techniques are also used by Ibanez (2004) and Carmona and Touzi (2008), where the latter uses the theory of the Snell envelope to determine the optimal exercise boundaries and also utilizes the Malliavin calculus for the computation of greeks. A constructive solution to the perpetual swing case for exponential Brownian motion is also given by Carmona and Touzi (2008). Unfortunately, these methods only work for the most basic versions of swing contracts where, at each time, only one unit of an option can be exercised.

More general swing contracts with a variable volume per exercise and an overall constraint can be priced with a tree-based method introduced by Jaillet *et al.* (2004).

In the above papers a discrete-time model for the underlying is used where one time step corresponds to the time frame in which no more than one right can be exercised, i.e. one day in most of the traded contracts. A special case where the number of exercise opportunities is equal to the number of exercise dates is considered by Howison and Rasmussen (2002) and a continuous optimal exercise strategy derived which yields a partial integro-differential equation for the option price.

Our method is based on the tree approach of Jaillet *et al.* (2004) with some slight modifications to adapt it to the peculiarities of our model for the underlying electricity price process.

27.4.1 The Grid Approach

The tree method of Jaillet *et al.* (2004) requires a discrete time model of the underlying. This is due to the fact that their swing contracts allow the holder to exercise at most one option within a specified time interval, say a day, and this is best modelled if the underlying process has the same time discretization. Assuming (S_t) is some continuous stochastic process for the spot price we obtain a discrete model by observing it on discrete points in time only, i.e.

$$S_{t_0}, S_{t_1}, S_{t_2}, \dots, S_{t_m},$$

with $t_0 = 0$, $t_{i+1} = t_i + \Delta t$, $t_m = T$ and $\Delta t = \frac{1}{365}$, indicating we can exercise on a daily basis.

Let the maturity date T be fixed and the pay-off at time t for simplicity* be given by $(S_t - K)^+$ for some strike price K and we assume only one unit of the underlying can be exercised in any time period. Let $V(n, s, t)$ denote the price of such a swing option at time t and spot price s which has n out of N exercise rights left. The dynamic programming principle allows us to write

* We could assume any general payoff function.

$$V(n, s, t) = \max_{n < N} \left\{ \begin{array}{l} e^{-r\Delta t} \mathbb{E}^Q[V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s], \\ e^{-r\Delta t} \mathbb{E}^Q[V(n-1, S_{t+\Delta t}, t + \Delta t) | S_t = s], \\ + (s - K)^+ \end{array} \right\}, \quad (27.16)$$

and $V(n, s, T) = (S - K)^+$, $0 < n \leq N$ and $V(0, s, t) = 0$. The conditional expectations can be written

$$\begin{aligned} & \mathbb{E}^Q[V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \\ &= \int_{-\infty}^{\infty} V(n, x, t + \Delta t) f_S(x; s) dx \end{aligned}$$

where $f_S(x; s)$ is the density of $S_{t+\Delta t}$ given $S_t = s$. Discretizing the spot variable we approximate

$$\begin{aligned} & \mathbb{E}^Q[V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \\ & \cong \sum_j V(n, s_j, t + \Delta t) f_S(s_j; s_i) \Delta s_j \end{aligned}$$

This is only one possible approximation; others might be to use higher-order integration rules or using only a few grid points in the sum based on the fact that $f_S(x; s) \rightarrow 0$ for $|s - x|$ large. For a trinomial tree, one only uses three grid points, i.e.

$$\mathbb{E}^Q[V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s] \approx \sum_{j=-1}^1 V(n, s_{i+j}, t + \Delta t) p_{i,j},$$

$p_{i,j}$ being the probability of going from node i to node j . However, such a tree approach is not well suited to our case for two reasons. First, the time step size is determined by the shortest time between two possible exercise dates, which is mainly one day for swing contracts. This limits the accuracy of the algorithm as a refinement of the grid in the spot direction will not improve the result. Second, in the presence of jumps, a three-point approximation for the conditional density is insufficient due to the heavy tails in the distribution. As a result, we keep our method general and say

$$\mathbb{E}^Q[V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s] \approx \sum_j V(n, s_j, t + \Delta t) p_{i,j},$$

where $p_{i,j}$ is an approximation to the density $f_S(s_j; s_i) \Delta s_j$ (it can accommodate higher-order integration rules and boundary approximations). With the notation $V_{i,k}^n := V(n, s_i, t_k)$ we can then write the method as

$$\begin{aligned} V_{i,k}^n &= \max \left\{ e^{-r\Delta t} \sum_j V_{j,k+1}^n p_{i,j}, e^{-r\Delta t} \sum_j V_{j,k+1}^{n-1} p_{i,j} + (s_i - K)^+ \right\}, \\ V_{i,k}^0 &= 0, \quad V_{i,m}^n = 0. \end{aligned} \quad (27.17)$$

27.4.2 Numerical Results

We now turn to the model of interest, (27.10), which exhibits spikes. Assume that the mean-reversion process (X_t) and the spike process (Y_t) are individually observable and so the value function V of a swing option depends on both variables and the general pricing principle (27.16) becomes

$$V(n, x, y, t) = \max \begin{cases} e^{-r\Delta t} \mathbb{E}^Q[V(n, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y], \\ e^{-r\Delta t} \mathbb{E}^Q[V(n-1, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y] + (e^{f(t)+x+y} - K)^+. \end{cases}$$

In order to calculate conditional expectations we need to define transition probabilities. Given one starts at node $(X_t, Y_t) = (x, y)$ the probability to arrive at node $(X_{t+\Delta t}, Y_{t+\Delta t}) = (x_k, y_l)$ is approximately given by

$$p_{t,j,k,t} \approx f_{X_{t+\Delta t}|X_t=x_i}(x_k) f_{Y_{t+\Delta t}|Y_t=y_i}(y_l) \Delta x \Delta y,$$

because X_t and Y_t are independent. The conditional density of the mean-reverting process (X_t) is known as $X_{t+\Delta t}$ given $X_t = x$ is normally distributed with $N(xe^{-\alpha\Delta t}, (\sigma^2/2\alpha)(1-e^{-2\alpha\Delta t}))$. As we do not have a closed-form expression for the density of the spike process we use approximations developed in [Section 27.2.1](#). For an exponential jump size distribution $J \sim \text{Exp}(1/\mu_j)$, for example, we use approximation (27.9) for the spike process at time t given zero initial conditions.

The introduction of a second space dimension increases the complexity of the algorithm considerably, essentially by a factor proportional to the square of the number of grid points in the y direction. To price the swing contract shown in Figure 27.10 which has 365 exercise dates and up to 100 exercise opportunities, our C++ implementation requires about 10 minutes to complete the calculation on an Intel P4, 3.4 GHz, and for a grid of 120×60 points in the x and y direction, respectively. The same computation but with no spikes and a grid of 120×1 points only takes about one second.

Based on Figure 27.10 we make two observations. First, the price per exercise right decreases with the number of exercise rights. This is the correct qualitative behaviour one would expect because n swing options each with one exercise right* only, offer more flexibility than one swing option with n exercise rights.† Second, the premium added due to the spike risk is much more significant for options with small numbers of exercise rights than for a large number. This is also intuitively clear, as an option with say 100 exercise rights will mainly be used against high prices caused by the diffusive part and only occasionally against spiky price explosions.




FIGURE 27.10 Value of a one year swing option per exercise right. Market parameters of the underlying are as before, see Figure 27.1, i.e. $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4$, $J \sim \text{Exp}(1/\mu_j)$ with $\mu_j = 0.4$, $f(t) = 0$, $r = 0$, and initial conditions $X_0 = 0$ and $Y_0 = 0$. The swing contract delivers over a time period of one year $T \in [0, 1]$ with up to 100 call rights and a strike price of $K = 1$, where a right can be exercised on any day. As a comparison the price of the same swing option is plotted but where the underlying does not exhibit spikes, i.e. $\lambda = 0$.

* This is actually an American option.

† The rights of a swing option can only be exercised one at a time.

In Figure 27.11 we show how sensitive swing option prices are to changes in market parameters. There we consider a swing option with a duration of 60 days and up to 20 exercise opportunities. In each graph we change one parameter by 20% up and down. The most significant change is caused by a change in the volatility parameter σ . Note, the long-term variance of the mean-reverting process (X_t) is $\sigma^2/2\alpha$ and we expect some direct relationship between the long-term variance and the option price. Hence, a change in the mean-reversion parameter α is inversely proportional to the price and quantitatively changes the price less than the volatility σ . The mean-reversion parameter β of the spike process has a similar effect on the option price as α has, but where the influence slightly decreases with the number of options. This is consistent with previous observations of the impact of jumps on option prices as seen in Figure 27.10. This effect is much more clearly visible for the other jump parameters λ and μ_j , which have the greatest impact on options with only a few exercise rights. For one exercise right, a 20% change in the jump size μ_j has an even greater effect on the price than a 20% change in volatility σ . A possible explanation is that we deal with the exponential of an exponentially distributed jump size which is very heavy tailed.




FIGURE 27.11 Sensitivity of swing option prices with respect to model parameters. A swing option with 60 exercise dates and up to 20 rights is considered, where the red curve is based on the parameters of Figure 27.10. In each graph, one market parameter is shifted up by 20% (green line) and down by 20% (blue line). We always plot the option price divided by the number of exercise rights.

References

- Benth, F., Ekeland, L., Hauge, R. and Neilsen, B., A note on arbitrage-free pricing of forward contracts in energy markets. *Appl. Math. Finan.*, 2003, **10**, 325–336.
- Benth, F., Kallsen, J. and Meyer-Brandis, T., A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivative pricing. *Appl. Math. Finan.*, 2007, **14**, 153–169.
- Black, F., The pricing of commodity contracts. *J. Finan. Econ.*, 1976, **3**, 167–179.
- Carmona, R. and Touzi, N., Optimal multiple stopping and valuation of swing options. *Math. Finan.*, 2008, **18**, 239–268.
- Carr, P. and Madan, D., Option valuation using the fast Fourier transform. *J. Comput. Finan.*, 1998, **2**, 61–73.
- Cartea, A. and Figueroa, M., Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Appl. Math. Finan.*, 2005, **12**, 313–335.
- Clewlow, L. and Strickland, C., *Energy Derivatives: Pricing and Risk Management*, 2000 (Lacima Publications: London).
- Cont, R. and Tankov, P., *Financial Modelling with Jump Processes*, 2004 (Chapman & Hall/CRC: New York).
- Deng, S., Stochastic models of energy commodity prices and their applications: mean-reversion with jumps and spikes. Working paper PWP-073, 2000.
- Duffie, D., Pan, J. and Singleton, K., Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 2000, **68**, 1343–1376.
- Howison, S. and Rasmussen, H., Continuous swing options. Working paper, 2002.
- Ibanez, A., Valuation by simulation of contingent claims with multiple early exercise opportunities. *Math. Finan.*, 2004, **14**, 223–248.
- Jaillet, P., Ronn, E. and Tompadis, S., Valuation of commodity based swing options. *Mgmt. Sci.*, 2004, **50**, 909–921.
- Kluge, T., Pricing swing options and other electricity derivatives. PhD thesis, University of Oxford, 2006.
- Klüppelberg, C., Meyer-Brandis, T. and Schmidt, A., Electricity spot price modelling with a view towards extreme spike risk. Preprint, 2008.
- Knittel, C.R. and Roberts, M.R., An empirical examination of restructured electricity prices. *Energy Econ.*, 2005, **27**, 791–817.
- Koopman, S.J., Ooms, M. and Carnero, M.A., Periodic seasonal Reg-ARFIMA-GARCH models for daily electricity spot prices. *J. Am. Statist. Assoc.*, 2007, **102**, 16–27.
- Lewis, A., A simple option formula for general jump diffusion and other exponential Lévy processes. Working paper, 2001.
- Lucia, J. and Schwartz, E., Electricity prices and power derivatives: evidence from the Nordic power exchange. *Rev. Deriv. Res.*, 2002, **5**, 5–50.
- Meinshausen, N. and Hambly, B., Monte Carlo methods for the valuation of options with multiple exercise opportunities. *Math. Finan.*, 2004, **14**, 557–583.
- Turnbull, S.M. and Wakeman, L.M., A quick algorithm for pricing European average options. *J. Finan. Quant. Anal.*, 1991, **26**, 377–389.
- Weron, R., Bierbrauer, M. and Trück, S., Modeling electricity prices: jump diffusion and regime switching. *Physica A*, 2004, **33**, 39–48.

28

Efficient Pricing of Swing Options in Lévy-Driven Models

28.1	Introduction	607
28.2	Lévy Processes: Basic Facts	608
	General Definitions • Regular Lévy Processes of Exponential Type	
	• The Wiener–Hopf Factorization	
28.3	The Multiple Optimal Stopping Problem for Swing Options	612
28.4	The Finite Difference Scheme for Pricing Swing Options	613
28.5	Pricing Swing Options Using the Wiener–Hopf Approach	615
28.6	Numerical Results	618
	Acknowledgements	619
	References	620

Oleg Kudryavtsev

Antonino Zanette

We consider the problem of pricing swing options with multiple exercise rights in Lévy-driven models. We propose an efficient Wiener–Hopf factorization method that solves multiple parabolic partial integro-differential equations associated with the pricing problem. We compare the proposed method with a finite difference algorithm. Both proposed deterministic methods are related to the dynamic programming principle and lead to the solution of a multiple optimal stopping problem. Numerical examples illustrate the efficiency and the precision of the proposed methods.

Keywords: Option pricing; Swing options; Finite difference methods; Wiener–Hopf factorization; American options; Energy derivatives; Numerical methods for option pricing

JEL Classification: C6, C61, C63, G1, G13

28.1 Introduction

The motivation for this work comes from energy markets, where financial instruments are becoming increasingly important for risk management. In a deregulated market, energy contracts will need to be priced according to their financial risk. Due to the uncertainty of consumption and the limited fungibility of energy, new financial contracts such as swing options have been introduced in the commodity market. Swing options are an American-type option with many exercise rights. Their owner can exercise them many times under the condition that they respect the refracting time that separates two successive exercises. Swing options are also widely used in the gas and oil markets. Thus, pricing swing options will become increasingly important.

Many numerical methods, essentially based on the solution to the dynamic programming equation, have been introduced recently in the financial literature. In the context of swing options, two different

probabilistic strategies have been developed. In the first, swing options are priced using an extension of the binomial tree algorithm, leading to the so-called forest tree (Lari-Lavassani *et al.* 2001, Jaillet *et al.* 2004). In the second, Monte Carlo methods are used, in which conditional expectations are computed using either regression techniques (Barrera-Esteve *et al.* 2006) or Malliavin calculus (Mnif and Zhegal 2006, Carmona and Touzi 2008). In particular, Carmona and Touzi (2008) propose a Monte Carlo approach to the problem of pricing American put options, in a finite time horizon, with multiple exercise rights in the case of geometric Brownian motion. In that paper, they introduce the inductive hierarchy of Snell envelopes needed in the multiple exercise case.

Energy expenditure increases sharply with higher daily temperature variation, and, consequently, the price varies. Although these spikes of power consumption are infrequent, they have a large financial impact, and therefore, many authors propose pricing swing options in a model with jumps (Mnif and Zhegal 2006, Wilhelm and Winter 2008). Mnif and Zhegal (2006) extend the results of Carmona and Touzi (2008) to a market with jumps. In fact, the multiple stopping time problem for swing options can be reduced to a cascade of single stopping time problems in a Lévy market where jumps are permitted. With regards to deterministic methods, Wilhelm and Winter (2008) develop a finite element algorithm for pricing swing options in different models, including jump models.

Boyarchenko and Levendorskii (2006) apply the Wiener–Hopf method to similar but distinct multistage investment or disinvestment problems (sequences of embedded perpetual American (real) options) under uncertainty, modelled as a monotone function of a Lévy process; in the case of the Kou model, closed-form solutions are given.

In this paper, we propose two approaches to solving multiple parabolic partial integro-differential equations (PIDEs) for pricing swing options in jump models. The first method, which is very simple and is introduced for comparison purposes, uses a finite difference scheme to solve the system of variational inequalities associated with the swing option problem approached through the splitting method proposed by Barles *et al.* (1995).

The second method uses fast Wiener–Hopf factorization (FWHF), introduced by Kudryavtsev and Levendorskii (2009), where a fast and accurate numerical method for pricing barrier options for a wide class of Lévy processes was constructed. The FWHF method is based on an efficient approximation of the Wiener–Hopf factors and the Fast Fourier Transform algorithm. The advantage of the Wiener–Hopf approach over finite difference schemes in terms of accuracy and convergence was demonstrated by Kudryavtsev and Levendorskii (2009). The FWHF-method was developed further in a series of papers by Kudryavtsev (2011, 2013). We will propose here a new efficient pricing algorithm for swing options that involves dynamic programming and the solving of multiple PIDEs by the FWHF method. We apply this algorithm for pricing swing options where the spot electricity price is a Lévy process that allows the consideration of jump risk. Numerical results, developed as in Wilhelm and Winter (2008) in the Black–Scholes and CGMY models, show the efficiency and accuracy of the proposed algorithms. The results presented in Chapter 28 were developed by Kudryavtsev and Zanette (2013).

The rest of the paper is organized as follows. Section 28.2 is devoted to the basic facts on Lévy processes. In Section 28.3, we present the multiple optimal stopping problem for swing options. In Sections 28.4 and 28.5, we propose, respectively, a finite difference and a Wiener–Hopf approach for pricing swing options. The numerical results are presented in Section 28.6.

28.2 Lévy Processes: Basic Facts

28.2.1 General Definitions

A Lévy process is a stochastically continuous process with stationary independent increments (for general definitions, see, e.g. Sato (1999)). A Lévy process may have a Gaussian component and/or a pure jump component. The latter is characterized by the density of jumps, which is called the Lévy density.

A Lévy process X_t can be completely specified by its characteristic exponent, ψ , definable from the equality $E[e^{i\xi X(t)}] = e^{-t\psi(\xi)}$ (we confine ourselves to the one-dimensional case).

The characteristic exponent is given by the Lévy–Khintchine formula:

$$\psi(\xi) = \frac{\sigma^2}{2} \xi^2 - i\mu\xi + \int_{-\infty}^{+\infty} (1 - e^{i\xi y} + i\xi y \mathbf{1}_{|y| \leq 1}) \nu(dy), \quad (28.1)$$

where $\sigma^2 \geq 0$ is the variance of the Gaussian component, and the Lévy measure $\nu(dy)$ satisfies

$$\int_{\mathbb{R} \setminus \{0\}} \min\{1, y^2\} \nu(dy) < +\infty \quad (28.2)$$

Assume that, under a risk-neutral measure chosen by the market, the price process has the dynamics $S_t = e^{X_t}$, where X_t is a certain Lévy process. Then we must have $E[e^{X_t}] < +\infty$, and, therefore, ψ must admit analytic continuation into a strip $\Im \xi \in (-1, 0)$ and continuous continuation into the closed strip $\Im \xi \in [-1, 0]$.

The infinitesimal generator of X , denoted by L , is an integro-differential operator that acts as follows:

$$\begin{aligned} Lu(x) = & \frac{\sigma^2}{2} \frac{\partial^2 u}{\partial x^2}(x) + \mu \frac{\partial u}{\partial x}(x) \\ & + \int_{-\infty}^{+\infty} \left(u(x+y) - u(x) - y \mathbf{1}_{|y| \leq 1} \frac{\partial u}{\partial x}(x) \right) \nu(dy). \end{aligned} \quad (28.3)$$

The infinitesimal generator L can also be represented as a pseudo-differential operator (PDO) with the symbol $-\psi(\xi)$, i.e. $L = -\psi(D)$, where $D = i\partial_x$. Recall that a PDO $A = a(D)$ acts as follows:

$$Au(x) = (2\pi)^{-1} \int_{-\infty}^{+\infty} e^{ix\xi} a(\xi) \hat{u}(\xi) d\xi, \quad (28.4)$$

where $\hat{u}$ is the Fourier transform of a function u ,

$$\hat{u}(\xi) = \int_{-\infty}^{+\infty} e^{-ix\xi} u(x) dx.$$

Note that the inverse Fourier transform in (28.4) is defined in the classical sense only if the symbol $a(\xi)$ and function $\hat{u}(\xi)$ are sufficiently nice. In general, one defines the (inverse) Fourier transform by duality.

Further, if the riskless rate, r , is constant, and if the stock does not pay dividends, then the discounted price process must be a martingale. Equivalently, the following condition (the EMM requirement) must hold (see, e.g. Boyarchenko and Levendorskii (2002)):

$$r + \psi(-i) = 0, \quad (28.5)$$

which can be used to express μ via the other parameters of the Lévy process:

$$\mu = r - \frac{\sigma^2}{2} + \int_{-\infty}^{+\infty} (1 - e^y + y \mathbf{1}_{|y| \leq 1}) \nu(dy). \quad (28.6)$$

Hence, the infinitesimal generator may be rewritten as follows:

$$\begin{aligned} Lu(x) = & \frac{\sigma^2}{2} \frac{\partial^2 u}{\partial x^2}(x) + \left(r - \frac{\sigma^2}{2} \right) \frac{\partial u}{\partial x}(x) \\ & + \int_{\mathbb{R}} \left[u(x+y) - u(x) - (e^y - 1) \frac{\partial u}{\partial x}(x) \right] \nu(dy). \end{aligned} \quad (28.7)$$

28.2.2 Regular Lévy processes of Exponential Type

Loosely speaking, a Lévy process X is called a *Regular Lévy Process of Exponential type* (RLPE) if its Lévy density has a polynomial singularity at the origin and decays exponentially at infinity (Boyarchenko and Levendorskiĭ 2002). An almost equivalent definition is as follows: the characteristic exponent is analytic in a strip $\Im \xi \in (\lambda_-, \lambda_+)$, $\lambda_- < -1 < 0 < \lambda_+$, is continuous up to the boundary of the strip, and admits the representation

$$\psi(\xi) = -i\mu\xi + \phi(\xi), \quad (28.8)$$

where $\phi(\xi)$ stabilizes to a positively homogeneous function at infinity:

$$\phi(\xi) \sim c_{\pm} |\xi|^{\nu}, \quad \text{as } \Re \xi \rightarrow \pm\infty, \text{ in the strip } \Im \xi \in (\lambda_-, \lambda_+), \quad (28.9)$$

where $c_{\pm} > 0$. ‘Almost’ means that the majority of classes of Lévy processes used in empirical studies of financial markets satisfy the conditions of both definitions. These classes are as follows: Brownian motion, Kou’s model (Kou 2002), Hyperbolic processes (Eberlein and Keller 1995), Normal Inverse Gaussian processes and their generalization (Barndorff-Nielsen 1998, Barndorff-Nielsen and Levendorskiĭ 2001), and the extended Koponen’s family. Koponen (1995) introduced a symmetric version; Boyarchenko and Levendorskiĭ (2000) gave a non-symmetric generalization; later, a subclass of this model appeared under the name the CGMY model in Carr *et al.* (2002), and Boyarchenko and Levendorskiĭ (2002) used the name KoBoL family.

The important exception is Variance Gamma Processes (VGP; see, e.g. Madan *et al.* (1998)). VGP satisfy the conditions of the first definition but not the second, as the characteristic exponent behaves like $\text{const} \cdot \ln |\xi|$ as $\xi \rightarrow \infty$.

Example 2.1

The characteristic exponent of a pure jump CGMY model is given by

$$\psi(\xi) = -i\mu\xi + C\Gamma(-Y)[G^Y - (G + i\xi)^Y + M^Y - (M - i\xi)^Y], \quad (28.10)$$

where $C > 0$, $\mu \in \mathbb{R}$, $Y \in (0, 2)$, $Y \neq 1$, and $-M < -1 < 0 < G$.

Example 2.2

If the Lévy measure of a jump diffusion process is given by a normal distribution

$$\nu(dx) = \frac{\lambda}{\delta\sqrt{2\pi}} \exp\left(-\frac{(x-\gamma)^2}{2\delta^2}\right) dx,$$

then we obtain the Merton model. The parameter λ characterizes the intensity of the jumps. The characteristic exponent of the process is of the form

$$\psi(\xi) = \frac{\sigma^2}{2}\xi^2 - i\mu\xi + \lambda\left(1 - \exp\left(-\frac{\delta^2\xi^2}{2} + i\gamma\xi\right)\right), \quad (28.11)$$

where $\sigma, \delta, \lambda \geq 0$, $\mu, \gamma \in \mathbb{R}$.

There are two important degenerate cases.

- If the intensity of jumps $\lambda = 0$, then we obtain the Black–Scholes model with $\mu = r - (\sigma^2/2)$ fixed by the EMM requirement.
- If the intensity of jumps $\lambda > 0$ but $\delta = 0$, then we obtain a jump diffusion process with a constant jump size γ ; the drift term $\mu = r - (\sigma^2/2) + \lambda(1 - e^\gamma)$ is fixed by the EMM requirement.

28.2.3 The Wiener–Hopf Factorization

There are several forms of the Wiener–Hopf factorization. The Wiener–Hopf factorization formula used in probability reads as follows:

$$E[e^{i\xi X_T}] = E[e^{i\xi \bar{X}_T}]E[e^{i\xi \underline{X}_T}], \quad \forall \xi \in \mathbb{R}, \quad (28.12)$$

where $T \sim \text{Exp } q$, and $\bar{X}_t = \sup_{0 \leq s \leq t} X_s$ and $\underline{X}_t = \inf_{0 \leq s \leq t} X_s$ are the supremum and infimum processes. Introducing the notation

$$\varphi_q^+(\xi) = qE\left[\int_0^\infty e^{-qt} e^{i\xi \bar{X}_t} dt\right] = E[e^{i\xi \bar{X}_T}], \quad (28.13)$$

$$\varphi_q^-(\xi) = qE\left[\int_0^\infty e^{-qt} e^{i\xi \underline{X}_t} dt\right] = E[e^{i\xi \underline{X}_T}], \quad (28.14)$$

we can write (28.12) as

$$\frac{q}{q + \psi(\xi)} = \varphi_q^+(\xi)\varphi_q^-(\xi). \quad (28.15)$$

Equation (28.15) is a special case of the Wiener–Hopf factorization of the symbol of a PDO. In applications to Lévy processes, the symbol is $q/(q + \psi(\xi))$, and the PDO is $\mathcal{E}_q := q/(q - L) = q(q + \psi(D))^{-1}$: the normalized resolvent of the process X_t or, using the terminology of Boyarchenko and Levendorskii (2005), the expected present value operator (EPV operator) of the process X_t . The name is due to the observation that, for a stream $g(X_t)$,

$$\mathcal{E}_q g(x) = E\left[\int_0^{+\infty} qe^{-qt} g(X_t) dt \mid X_0 = x\right].$$

We introduce the following operators:

$$\mathcal{E}_q^\pm := \varphi_q^\pm(D), \quad (28.16)$$

which also admit interpretation as the EPV operators under supremum and infimum processes. One of the basic observations in the theory of PDO is that the product of symbols corresponds to the product of operators. In our case, it follows from (28.15) that

$$\mathcal{E}_q = \mathcal{E}_q^+ \mathcal{E}_q^- = \mathcal{E}_q^- \mathcal{E}_q^+ \quad (28.17)$$

as operators in appropriate function spaces.

For a wide class of Lévy models, $\mathcal{E}$ and $\mathcal{E}^\pm$ admit interpretation as expectation operators:

$$\begin{aligned} \mathcal{E}_q g(x) &= \int_{-\infty}^{+\infty} g(x+y) P_q(y) dy, \\ \mathcal{E}_q^\pm g(x) &= \int_{-\infty}^{+\infty} g(x+y) P_q^\pm(y) dy, \end{aligned}$$

where $P_q(y)$ and $P_q^\pm(y)$ are certain probability densities with

$$P_q^\pm(y) = 0, \quad \forall \pm y < 0.$$

Moreover, the characteristic functions of the distributions $P_q(y)$ and $P_q^\pm(y)$ are $q(q + \psi(\xi))^{-1}$ and $\varphi_q^\pm(\xi)$, respectively.

The general results in this paper are based on simple properties of the EPV operators, which follow immediately from the interpretation of $\mathcal{E}^\pm$ as expectation operators. For details, see Boyarchenko and Levendorskiĭ (2005).

Proposition 2.3

EPV operators $\mathcal{E}_q^\pm$ have the following properties.

- (1) If $g(x) = 0 \quad \forall x \geq h$, then $\forall x \geq h$, $(\mathcal{E}_q^+ g)(x) = 0$ and $((\mathcal{E}_q^+)^{-1} g)(x) = 0$.
- (2) If $g(x) = 0 \quad \forall x \geq h$, then $\forall x \geq h$, $(\mathcal{E}_q^- g)(x) = 0$ and $((\mathcal{E}_q^-)^{-1} g)(x) = 0$.
- (3) If $g(x) \geq 0 \quad \forall x$, then $(\mathcal{E}_q^+ g)(x) \geq 0 \quad \forall x$. If, in addition, there exists x_0 such that $g(x) > 0 \quad \forall x > x_0$, then $(\mathcal{E}_q^+ g)(x) > 0 \quad \forall x$.
- (4) If $g(x) \geq 0 \quad \forall x$, then $(\mathcal{E}_q^- g)(x) \geq 0 \quad \forall x$. If, in addition, there exists x_0 such that $g(x) > 0 \quad \forall x < x_0$, then $(\mathcal{E}_q^- g)(x) > 0 \quad \forall x$.
- (5) If g is monotone, then $\mathcal{E}_q^+ g$ and $\mathcal{E}_q^- g$ are also monotone.
- (6) If g is continuous and satisfies

$$|g(x)| \leq C(e^{\sigma_- x} + e^{\sigma_+ x}), \quad \forall x \in \mathbb{R}, \quad (28.18)$$

where $\sigma_- \geq 0 \geq \sigma_+$ and C are all independent of x , then $\mathcal{E}_q^+ g$ and $\mathcal{E}_q^- g$ are continuous.

28.3 The Multiple Optimal Stopping Problem for Swing Options

We consider a price process that evolves according to the formula

$$S_t = e^{X_t},$$

where $\{X_t\}_{t \geq 0}$, the driving process, is an adapted Lévy process defined on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, satisfying the usual conditions.

Let T be the option's maturity time, and let $\mathcal{T}_{t,T}$ be the set of $\mathbb{F}$ -stopping times with values in $[t, T]$. Consider a swing option that gives the right to multiple exercise with a $\delta > 0$ refracting period that separates two successive exercises (the number of possible exercises is fixed). We consider the possibility of n put exercises. We shall denote by $\mathcal{T}^n$ the collection of all vectors of stopping times $(\tau_1, \tau_2, \dots, \tau_n)$, such that

- $\tau_1 \geq T$ a.s.;
- $\tau_i - \tau_{i-1} \geq \delta$ on $\{\tau_{i-1} \leq T\}$ a.s. for all $i = 2, \dots, n$.

Denote by $\nu^{(i)}(t, x)$ the swing option value with the possibility of i exercises at spot level $S = e^x$ and time $t \geq T$.

Following Carmona and Touzi (2008), the multiple exercise problem can be solved by computing

$$\nu^{(n)}(0, x) = \sup_{(\tau_1, \dots, \tau_n) \in \mathcal{T}^n} \sum_{i=1}^n E[e^{-r\tau_i} \phi(X_{\tau_i})], \quad (28.19)$$

where

$$\phi(x) = (K - e^x)_+$$

is the pay-off function.

To solve the multiple optimal stopping problem, Carmona and Touzi (2008) introduce the idea of an inductive hierarchy. In fact, they reduce the multiple stopping problem to a cascade of n optimal single stopping problems. Define the value function for $i = 1, \dots, n$

$$v^{(i)}(t, x) = \sup_{\tau \in \mathcal{T}_{t,T}} E[e^{-r\tau} \phi^{(i)}(\tau, X_{\tau}^{t,x})], \quad (28.20)$$

where the reward function $\phi^{(i)}$ is now defined as

$$\phi^{(i)}(t, x) = \phi(x) + E[e^{-r\delta} v^{(i-1)}(t + \delta, X_{t+\delta}^{t,x})], \quad t \leq T - \delta, \quad (28.21)$$

$$\phi^{(i)}(t, x) = \phi(x), \quad t > T - \delta. \quad (28.22)$$

The problem can be solved using a Monte Carlo algorithm. Let $t_0 = 0 < t_1 < t_2 < \dots < t_N = T$ be a time discretization grid. The price of a swing option can be computed by the backward induction procedure

$$\begin{cases} v^{(i)}(t_N, x) = \phi(x), \\ v^{(i)}(t_{k-1}, x) = \max\{\phi^{(i)}(t_{k-1}, x); \\ e^{-r(t_k - t_{k-1})} E[v^{(i)}(t_k, X_{t_k}^{t_{k-1}, x})]\}, \quad k = N, \dots, 1. \end{cases}$$

Carmona and Touzi (2008) and Mnif and Zeghal (2006), respectively, considered a Monte Carlo Malliavin-based algorithm to compute the price in the Black–Scholes and jump models frameworks. Barrera-Esteve *et al.* (2006) used a regression-based method to approximate conditional expectations. In the following sections, we propose two PIDE-based approaches.

28.4 The Finite Difference Scheme for Pricing Swing Options

We can compute the swing option price using the formulation given in (28.20) with an analytical approach. In fact, we propose to solve the following system of variational inequalities associated with the swing options formulation:

$$\begin{cases} \max\left(\phi^{(i)}(t, x) - v^{(i)}(t, x), \frac{\partial v^{(i)}}{\partial t} + Lv^{(i)} - rv^{(i)}\right) = 0 \\ (t, x) \in [0, T] \times \mathbb{R}, \\ v^{(i)}(T, x) = \phi^{(i)}(T, e^x), \end{cases} \quad (28.23)$$

with $i = 1, \dots, n$, where the integro-differential operator L is defined in (28.7).

Now recall that, for $t \geq T - \delta$,

$$\phi^{(i)}(t, x) = (K - e^x)_+ + E[e^{-r\delta} v^{(i-1)}(t + \delta, X_{t+\delta}^{t,x})].$$

Let us define for $t \leq T - \delta$

$$u^{(i)}(t, x) = E[e^{-r\delta} v^{(i)}(t + \delta, X_{t+\delta}^{t,x})].$$

By the Feynman–Kac theorem, $u^{(i)}(t, x) = z(0, x)$, where $z(t, x)$ is the solution of the following partial integro-differential equation (PIDE):

$$\begin{cases} \frac{\partial z}{\partial t} + Lz - rz = 0, & (t, x) \in [0, \delta] \times \mathbb{R}, \\ z(\delta, x) = v^{(i)}(t + \delta, x), \end{cases} \quad (28.24)$$

which can be computed numerically using a finite difference approach. To price a swing option, therefore, we can solve the system of variational inequalities (28.23) computing the reward pay-off function $\phi^{(i)}(t, x)$ in the following way:

$$\phi^{(i)}(t, x) = \phi(x).$$

for $T - \delta < t \leq T$, and

$$\phi^{(i)}(t, x) = (K - e^x)_+ + u^{(i-1)}(t, x),$$

for $t \leq T - \delta$. As stated above, the reward pay-off function can be computed numerically using a finite difference scheme. The numerical solution of the variational inequalities (28.23) requires numerically solving each PIDE problem (28.24). To solve (28.23) and (28.24), we perform the following steps.

- *Localization.* We choose a spatial bounded computational domain Ω_T , which implies that we must choose some artificial boundary conditions.
- *Truncation of large jumps.* This step corresponds to truncating the integration domain in the integral part.
- *Discretization.* The derivatives of the solution are replaced by finite differences, and the integral terms are approximated using the trapezoidal rule. Then the problem is solved by using an explicit–implicit scheme (see Briani *et al.* (2004) and Cont and Voltchkova (2005) and the program implementation in the PREMIA software (www.premia.fr)). In particular, we introduce a time grid $t = s\Delta t$, $s = 0, \dots, N$, where $\Delta t = T/N$ is the time step. At each time step, it is necessary to solve a linear system for the linear problem (28.24) and a linear complementarity problem for the nonlinear problem (28.23). The idea of the explicit–implicit method is based on an asymmetric treatment of the differential and integral parts of L . The operator L in (28.24) is split into two parts,

$$Lz = Dz + Jz,$$

- where D and J are the differential and integral parts of L , respectively. We replace Dz with a finite difference approximation $D_\Delta z$ and Jz with the trapezoidal quadrature approximation $J_\Delta z$ and use the following explicit–implicit time-stepping:

$$\frac{z^{s+1} - z^s}{\Delta t} + D_\Delta z^s + J_\Delta z^{s+1} - rz^s = 0.$$

The integral part is treated in an explicit way to avoid a dense matrix, while the differential part is treated in an implicit way. Details of the algorithms are given in Cont and Voltchkova (2005).

- *Treatment of the variational inequalities.* We solve each of the variational inequalities (28.23) using the splitting method of Barles *et al.* (1995). The splitting methods can be viewed as an analytical version of dynamic programming. The idea of this scheme is to split the American problem

into two steps: we construct recursively the approximate solution $v^{(i)}(s\Delta t, x)$ at each time step $s\Delta t$ by starting from $v^{(i)}(N\Delta t, x) = \phi(x)$ and computing at each time step the values of $v^{(i)}(s\Delta t, x)$ for $s = N - 1, \dots, 0$ as follows:

- compute the solution of the following linear Cauchy problem on $[s\Delta t, (s+1)\Delta t] \times \Omega_t$ using an explicit–implicit scheme:

$$\begin{cases} \frac{\partial w^{(i)}(s\Delta t, x)}{\partial t} + Lw^{(i)}(s\Delta t, x) - rw^{(i)}(s\Delta t, x) = 0 \\ \text{in } [s\Delta t, (s+1)\Delta t] \times \Omega_t, \\ w^{(i)}((s+1)\Delta t, x) = v^{(i)}((s+1)\Delta t, x); \end{cases}$$

- apply the early exercise $v^{(i)}(s\Delta t, x) = \max(w^{(i)}(s\Delta t, x), \phi^{(i)}(s\Delta t, x))$, where the reward function $\phi^{(i)}(s\Delta t, x)$ is obtained by solving the linear problem (28.24) with an explicit–implicit finite difference method.

One could also apply the method of horizontal lines or Carr's randomization to (28.20), then use the explicit–implicit finite difference scheme to solve the corresponding sequence of free boundary problems. The analytical method of lines was introduced to finance by Carr and Faguet (1994); Carr (1998) suggested an important new probability interpretation of the method, which we call Carr's randomization. In the case of American options, the convergence of Carr's randomization algorithm was proved by Bouchard *et al.* (2005) for a wide class of strong Markov processes. In the next section, we will start with Carr's randomization procedure.

28.5 Pricing Swing Options Using the Wiener–Hopf Approach

In this section, we apply the Wiener–Hopf approach to pricing swing options. The first step is to discretize the time $(0 =) t_0 < t_1 < \dots < t_N (= T)$, but not the space variable. Set $v_N^i(x) = (K - e^x)_+$. For $s = N - 1, N - 2, \dots, 0$, set $\Delta_s = t_{s+1} - t_s$, $q^s = r + (\Delta_s)^{-1}$, and denote by $v_s^i(x)$ Carr's randomized approximation to $v^i(t_s, x)$.

The early exercise boundary h_s^i for an interval (t_s, t_{s+1}) and $v_s^i(x)$ can be found using backward induction. For $s = N - 1, N - 2, \dots$, the boundary h_s^i is chosen to maximize

$$v_s^i(x) = E \left[\int_0^{\tau_s^i} e^{-q^s t} v_{s+1}^i(X_t^{0,x}) dt \right] + E[e^{-q^s \tau_s^i} \phi_s(X_{\tau_s^i}^{0,x})], \quad (28.25)$$

where τ_s^i is the hitting time of the interval of the form $(-\infty, h_s^i]$, and

$$\phi_s^{(i)}(x) = (K - e^x) + E[e^{-r\delta} v^{(i-1)}(t_s + \delta, X_{t_s+\delta}^{t_s, x})], \quad t_s \leq T - \delta,$$

and

$$\phi_s^{(i)}(x) = (K - e^x), \quad t_s > T - \delta.$$

As in Boyarchenko and Levendorskii (2009), where the case of American options was considered, to derive (28.25), we replace $\phi(x) = (K - e^x)_+$ in (28.19) with $(K - e^x)$. This replacement is justified by a simple consideration that it is non-optimal to exercise the option when $(K - e^x) \leq 0$.

In this paper, we use uniform spacing; therefore, q^s and Δ_s are independent of s and denoted q and Δt , respectively. For the case of put swing options, v_s^i given by (28.25) is a unique solution of the boundary problem

$$(q-L)v_s^i(x) = (\Delta t)^{-1}v_{s+1}^i(x), \quad x > h_s^i, \quad (28.26)$$

$$v_s^i(x) = \phi_s^{(i)}(x), \quad x \leq h_s^i. \quad (28.27)$$

Note that the problem (28.26)–(28.27) can be obtained by discretization of the time derivative in the generalized Black–Scholes equation (see details in Boyarchenko and Levendorskii (2009) and the bibliography therein).

Let the refracting period δ be equal to $k\Delta t$, where k is a certain positive integer. Then, for $i = 1, \dots, n$,

$$\phi_s^{(i)}(x) = (K - e^x) + u_s^{i-1}(x), \quad (28.28)$$

where

$$u_s^0(x) = 0, \quad (28.29)$$

$$u_s^i(x) = 0, \quad t_s > T - \delta, \quad (28.30)$$

$$u_s^i(x) = E \left[e^{-r\delta} v_{s+k}^{(i)} \left(X_{t_s+k}^{t_s, x} \right) \right], \quad t_s \leq T - \delta. \quad (28.31)$$

Introduce $\tilde{v}_s^i(x) = v_s^i(x) - \phi_s^{(i)}(x)$, and substitute $v_s^i(x) = \tilde{v}_s^i(x) + \phi_s^{(i)}(x)$ into (28.26)–(28.27) as follows:

$$(q-L)\tilde{v}_s^i(x) = (\Delta t)^{-1}G_s^i(x), \quad x > h_s^i, \quad (28.32)$$

$$\tilde{v}_s^i(x) = 0, \quad x \leq h_s^i, \quad (28.33)$$

where $G_s^i = \tilde{v}_{s+1}^i + \phi_{s+1}^{(i)} - \Delta t(q-L)\phi_s^{(i)}$.

Using arguments similar to those of Boyarchenko and Levendorskii (2009), it can be shown that, for $s = n-1, n-2, \dots, 0$, the function G_s^i is a non-decreasing continuous function satisfying bound (28.18) with $\sigma_+ = 1, \sigma_- = 0$; in addition,

$$G_s^i(-\infty) < 0 < G_s^i(+\infty) = +\infty. \quad (28.34)$$

Then $G_s^i(x)$ satisfies the conditions of theorem 2.6 of Boyarchenko and Levendorskii (2009). Due to this theorem and proposition 2.3, the following statements hold.

1. The function

$$\tilde{w}_s^i := \varepsilon_q^+ G_s^i \quad (28.35)$$

is continuous; it increases and satisfies (28.34).

2. The equation

$$\tilde{w}_s^i(h) = 0 \quad (28.36)$$

has a unique solution, denoted by h_s^i .

3. The hitting time of $(-\infty, h_s^i]$, $\tau(h_s^i)$, is a unique optimal stopping time.
4. (Carr's approximation to) the swing option value with i exercise rights at the moment s is given by

$$v_s^i = (q\Delta t)^{-1} \varepsilon_q^- \mathbf{1}_{(h_s^i, +\infty)} \tilde{w}_s^i + \phi_s^{(i)}. \quad (28.37)$$

Equivalently,

$$\tilde{v}_s^i = (q\Delta t)^{-1} \varepsilon_q^- \mathbf{1}_{(h_s^i, +\infty)} \tilde{w}_s^i. \quad (28.38)$$

5. $\tilde{v}_s^i = v_s^i - \phi_s^{(i)}$ is a positive non-decreasing function that admits bound (28.18) with $\sigma_+ = 1$, $\sigma_- = 0$ and satisfies $\tilde{v}_s^i(+\infty) = +\infty$; it vanishes below h_s^i and increases on $[h_s^i, +\infty)$.

Because functions G_s^i and $\tilde{w}_s^i$ tend to plus infinity as $x \rightarrow +\infty$, the numerical calculation of the integrals in (28.35) and (28.37) may face certain difficulties. To improve the convergence, we reformulate the algorithm in terms of the bounded functions v_s^i . Taking into account (28.28) and (28.5), G_s^i can be rewritten as follows:

$$\begin{aligned} G_s^i(x) &= v_{s+1}^i(x) - \Delta t(q-L)u_s^{(i-1)}(x) - \Delta t(q-L)(K - e^x) \\ &= v_{s+1}^i(x) - \tilde{u}_s^{(i-1)}(x) - (\Delta t K q - e^x), \end{aligned} \quad (28.39)$$

where $\tilde{u}_s^i(x)$ can be approximated by the formulae

$$\tilde{u}_s^0(x) = 0, \quad (28.40)$$

$$\tilde{u}_s^i(x) = 0, \quad t_s > T - \delta, \quad (28.41)$$

$$\tilde{u}_s^i(x) = E[e^{-r(\delta-\Delta t)} \nu_{s+k}^{(i)}(X_{t_{s+k}}^{t_{s+1}, x})] + o(\Delta t), \quad t_s \leq T - \delta. \quad (28.42)$$

Note that we can easily compute the expectation on the RHS of (28.42) using the Fourier transform technique (see, e.g. Carr and Madan (1999) or Boyarchenko and Levendorskiĭ (2002)) as follows:

$$\tilde{u}_s^i(x) \approx (2\pi)^{-1} e^{-\rho x} \int_{-\infty}^{+\infty} e^{ix\xi - (\delta-\Delta t)(r+\psi(\xi+ip))} \hat{\nu}_{s+k}^{(i),\rho}(\xi) d\xi, \quad (28.43)$$

where $\hat{\nu}_s^{(i),\rho}(\xi)$, is the Fourier transform of the price $\nu_s^{(i)}(x)$ multiplied by an appropriate damping exponential factor $e^{\rho x}$; in our case, $\rho > 0$. Numerically, formula (28.43) can be efficiently realized by means of the FFT technique (Carr and Madan 1999).

However, for very short refracting periods δ , the integrand in (28.43) may decay slowly at infinity (see, e.g. Lord *et al.* (2008)). Hence, the numerical implementation of the Fourier transform may not be sufficiently accurate. To circumvent the potential numerical pricing difficulties when dealing with the case $\delta = k\Delta t$, $k > 1$, (k is not too large and Δt is small), the finite difference approach proposed in Section 28.4 can be used efficiently to find $\tilde{u}_s^i(x)$. The expectation in (28.42) can be interpreted as the solution at time Δt to the problem (28.24), with t_s instead of t . Finally, if $\delta = \Delta t$, then $\tilde{u}_s^i(x) \approx \nu_{s+k}^{(i)}(x)$.

We can now rewrite (28.37) as follows:

$$v_s^i = (q\Delta t)^{-1} \varepsilon_q^- (\mathbf{1}_{(h_s^i, +\infty)} w_s^i - \mathbf{1}_{(-\infty, h_s^i]} w_{s,0}^i), \quad (28.44)$$

where

$$w_s^i = \varepsilon_q^+ v_{s+1}^i, \quad (28.45)$$

$$\begin{aligned} w_{s,0}^i &= \varepsilon_q^+ (\tilde{u}_s^{(i-1)}(x) + \Delta t K q - e^x) \\ &= \varepsilon_q^+ \tilde{u}_s^{(i-1)}(x) + \Delta t K q - \varphi_q^+(-i) e^x, \end{aligned} \quad (28.46)$$

and h_s^i is a solution to the equation

$$w_s^i = w_{s,0}^i. \quad (28.47)$$

Note that, in (28.44) and (28.45), the functions in the arguments of the operators ε_q^- and ε_q^+ are bounded. The algorithm can be efficiently realized by using the Fast Wiener–Hopf factorization method (see details in Section 2 of Kudryavtsev and Levendorskii (2009)).

28.6 Numerical Results

In this section, we illustrate the efficiency and the robustness of the proposed methods numerically using the parameters of the numerical examples for pricing swing options in the Black–Scholes and CGMY models provided by Wilhelm and Winter (2008).

We consider a put swing option with $n = 1, 2, 3$ exercise numbers and a refracting period $\delta = 0.1$. We assume that the initial value of the stock prices is $S = 100$, the exercise price $K = 100$, the maturity $T = 1$ and the force of the interest rate $r = 0.05$.

In order to solve the PIDE numerically using the finite difference scheme, we first localize the variables and the integral term to bounded domains. We use for this purpose the estimates for the localization domain and the truncation of large jumps given by Voltchkova and Tankov (2008).

In the case of the Wiener–Hopf approach, we use the adaptive method from Kudryavtsev and Levendorskii (2009). For a fixed number of time steps, N , and a step in x -space, Δx , we increase the domain in x -space twofold to ensure that the price does not change significantly. In the dual space it corresponds to increasing the number of points M . Fix the space step $\Delta x > 0$ and the number of space points (and dual space points) $M = 2^m$. Define the partitions of the normalized log-price domain $[-M\Delta x/2; M\Delta x/2]$ by points $x_k = -(M\Delta x/2) + k\Delta x$, $k = 0, \dots, M-1$, and the frequency domain $[-\pi/\Delta x; \pi/\Delta x]$ by points $\xi_l = 2\pi l/\Delta x$, $l = -M/2, \dots, M/2$.

In our examples, both methods use a spatial discretization step $\Delta x = 0.001$ and a varying number of time steps $N = 50, 100, 200$. In the Wiener–Hopf approach, the optimal choice for the number of space points is $M = 2^{12}$ (doubling the number M changes the option prices by 0.0001% or less).

We propose first to assess the numerical robustness of our algorithm in the Black–Scholes case, using the volatility $\sigma = 0.3$. Table 28.1 reports the prices (with time in seconds in parentheses) with the relative errors in a Black–Scholes framework, using the finite difference method (FD) proposed in Section 28.4 and the Wiener–Hopf approach (FWHF) proposed in Section 28.5. As benchmark solutions, we take those provided by Wilhelm and Winter (2008) (B-WW).

Furthermore, we provide numerical results in a Lévy market model. To be exact, we use the CGMY model (Carr *et al.* 2002) with $C = 1$, $G = 10$, $M = 10$ and $Y = 0.5$. No comparison results are available in the paper of Wilhelm and Winter (2008), and thus we use as the benchmark value the FWHF method with a very fine mesh grid ($\Delta x = 0.0002$, $N = 800$ and $M = 2^{15}$). In order to justify the ‘benchmark’ values, we observed that with a finer mesh grid ($\Delta x = 0.0001$, $N = 1600$ and $M = 2^{17}$) the option prices change only by 0.02% or less. Table 28.2 reports the numerical results for the CGMY model.

Table 28.3 reports the prices of swing options in the CGMY model with decreasing values of refracting periods $\delta = 0.1, 0.01, 0.001, 0$. The prices are calculated using the FWHF method with a spatial

TABLE 28.1 Swing Option Prices in the Black–Scholes Model

		Prices			Relative errors (%)		
	N	$n = 1$	$n = 2$	$n = 3$	$n = 1$	$n = 2$	$n = 3$
FWHF	50	9.786 (0.22)	19.130 (0.7)	27.968 (1.17)	−0.85	−0.65	−0.56
	100	9.826 (0.42)	19.190 (1.38)	28.045 (2.33)	−0.45	−0.34	−0.29
	200	9.848 (0.83)	19.222 (2.72)	28.085 (4.59)	−0.22	−0.17	−0.15
B-WW		9.8700	19.2550	28.1265			
FD	50	9.834 (0.12)	19.096 (5.16)	27.711 (8.10)	−0.36	−0.83	−1.48
	100	9.864 (0.27)	19.184 (8.64)	27.925 (16.47)	−0.06	−0.37	−0.72
	200	9.867 (0.57)	19.220 (16.68)	28.027 (33.0)	−0.03	−0.18	−0.35

TABLE 28.2 Swing Option Prices in the CGMY Model

		Prices			Relative errors (%)		
	N	$n = 1$	$n = 2$	$n = 3$	$n = 1$	$n = 2$	$n = 3$
FWHF	50	7.100 (0.22)	13.859 (0.7)	20.228 (1.17)	−0.80	−0.61	−0.54
	100	7.131 (0.42)	13.905 (1.38)	20.287 (2.33)	−0.36	−0.28	−0.25
	200	7.147 (0.83)	13.928 (2.72)	20.317 (4.59)	−0.14	−0.11	−0.10
B-FWHF		7.157	13.944	20.337			
FD	50	7.173 (1.20)	13.887 (37.1)	20.102 (76.1)	0.22	−0.41	−1.16
	100	7.172 (2.31)	13.928 (146)	20.238 (286.5)	0.21	−0.11	−0.49
	200	7.171 (4.56)	13.948 (751)	20.306 (1398)	0.20	0.03	−0.15

TABLE 28.3 Convergence of Swing Option Prices in the CGMY Model

	δ	$n = 3$
FWHF	0.1	20.3416
	0.01	21.3701
	0.001	21.4699
	0	21.4760

discretization step $\Delta x = 0.001$ and number of time steps $N = 1000$; the other parameters remain the same. We see that the sequence of prices increases to the limit value as the refracting period goes to 0. The limit value is the solution to problem (28.20) with the reward function in (28.21) defined for zero refracting period. Hence, in the limit case, we obtain a simplified problem, because we do not need to calculate the expectation in (28.21).

All computations were performed in double precision on an Eee PC with the following characteristics: CPU Atom N450, 1.67 GHz, 2 Gb of RAM.

The numerical results confirm the reliability of both approaches, demonstrating the robustness of the methods. In particular, the Wiener–Hopf approach is undoubtedly a very precise and efficient method for pricing swing options in the presence of multiple jumps.

Acknowledgements

The first author gratefully acknowledges financial support from the European Science Foundation (ESF) through the Short Visit Grant number 3404 of the program ‘Advanced Mathematical Methods for Finance’ (AMaMeF) and the Russian Foundation for Humanities, the project number 15–32–01390.

References

- Barles, G., Daher, C. and Romano, M., Convergence of numerical schemes for problems arising in finance theory. *Math. Models Meth. Appl. Sci.*, 1995, **5**(1), 125–143.
- Barndorff-Nielsen, O.E., Processes of normal inverse Gaussian type. *Finance Stochast.*, 1998, **2**, 41–68.
- Barndorff-Nielsen, O.E. and Levendorskiĭ, S., Feller processes of normal inverse Gaussian type. *Quant. Finance*, 2001, **1**, 318–331.
- Barrera-Esteve, C., Bergeret, F., Dossal, C., Gobet, E., Meziou, A., Munos, R. and Reboul-Salze, D., Numerical methods for the pricing of swing options: A stochastic control approach. *Methodol. Comput. Appl. Probab.*, 2006, **8**(4), 517–540.
- Bouchard, B., El Karoui, N. and Touzi, N., Maturity randomization for stochastic control problems. *Ann. Appl. Probab.*, 2005, **15**(4), 2575–2605.
- Boyarchenko, S.I. and Levendorskiĭ, S.Z., Option pricing for truncated Lévy processes. *Int. J. Theor. Appl. Finance*, 2000, **3**, 549–552.
- Boyarchenko, S.I. and Levendorskiĭ, S.Z., *Non-Gaussian Merton–Black–Scholes Theory*, 2002 (World Scientific: Singapore).
- Boyarchenko, S.I. and Levendorskiĭ, S.Z., American options: The EPV pricing model. *Ann. Finance*, 2005, **1**(3), 267–292.
- Boyarchenko, S.I. and Levendorskiĭ, S.Z., General option exercise rules, with applications to embedded options and monopolistic expansion. *Contrib. Theor. Econ.*, 2006, **6**(1), article 2.
- Boyarchenko, S.I. and Levendorskiĭ, S.Z., Pricing American options in regime-switching models. *SIAM J. Control Optimiz.*, 2009, **48**, 1353–1376.
- Briani, M., La Chioma, C. and Natalini, R., Convergence of numerical schemes for viscosity solutions to integro-differential degenerate parabolic problems arising in financial theory. *Numer. Math.*, 2004, **98**(4), 607–646.
- Carmona, C. and Touzi, N., Optimal multiple stopping and valuation of swing options. *Math. Finance*, 2008, **18**(2), 239–268.
- Carr, P., Randomization and the American put. *Rev. Financ. Stud.*, 1998, **11**, 597–626.
- Carr, P. and Faguet, D., Fast accurate valuation of American options. Working Paper, Cornell University, 1994.
- Carr, P., Geman, H., Madan, D.B. and Yor, M., The fine structure of assets return: An empirical investigation. *J. Business*, 2002, **75**(3), 305–332.
- Carr, P. and Madan, D., Option valuation using the Fast Fourier Transform. *J. Comput. Finance*, 1999, **3**, 463–520.
- Cont, R. and Voltchkova, E., A finite difference scheme for option pricing in jump-diffusion and exponential Lévy models. *SIAM J. Numer. Anal.*, 2005, **43**(4), 1596–1626.
- Eberlein, E. and Keller, U., Hyperbolic distributions in finance. *Bernoulli*, 1995, **1**, 281–299.
- Jaillet, P., Ronn, E.I. and Tompaidis, S., Valuation of commodity-based swing options. *Mgmt Sci.*, 2004, **50**, 909–921.
- Koponen, I., Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process. *Phys. Rev. E*, 1995, **52**, 1197–1199.
- Kou, S.G., A jump-diffusion model for option pricing. *Mgmt Sci.*, 2002, **48**, 1086–1101.
- Kudryavtsev, O.E. and Levendorskiĭ, S.Z., Fast and accurate pricing of barrier options under Levy processes. *Finance Stochast.*, 2009, **13**(4), 531–562.
- Kudryavtsev, O.Ye., An efficient numerical method to solve a special case of integro-differential equations relating to the Levy models. *Mathematical Models and Computer Simulations*, 2011, **3**(6), 706–711.
- Kudryavtsev, O., Finite difference methods for option pricing under Levy processes: Wiener–Hopf factorization approach. *The Scientific World Journal*, 2013, **2013**, Article ID 963625, 12 pages
- Kudryavtsev, O. and Zanette, A., Efficient pricing of swing options in Levy-driven models. *Quantitative Finance*, 2013, **13**(4), 627–635.

- Lari-Lavassani, A., Simchi, M. and Ware, A., A discrete valuation of swing options. *Can. Appl. Math. Q.*, 2001, **9**(1), 35–73.
- Lord, R., Fang, F., Bervoets, F. and Oosterlee, C.W., A fast and accurate FFT-based method for pricing early-exercise options under Levy processes. *SIAM J. Sci. Comput.*, 2008, **30**(4), 1678–1705.
- Madan, D.B., Carr, P. and Chang, E.C., The variance Gamma process and option pricing. *Eur. Finance Rev.*, 1998, **2**, 79–105.
- Mnif, M. and Zhegal, A.B., Optimal multiple stopping and valuation of swing options in Lévy models. *Int. J. Theor. Appl. Finance*, 2006, **9**(8), 1267–1298.
- Sato, K., *Lévy Processes and Infinitely Divisible Distributions*, 1999 (Cambridge University Press: Cambridge).
- Voltchkova, E. and Tankov, P., Deterministic methods for option pricing in exponential Lévy models. *PREMIA documentation*, 2008. Available online at: <http://www.premia.fr>
- Wilhelm, M. and Winter, C., Finite element valuation of swing options. *J. Comput. Finance*, 2008, **11**(3), 107–132.

Hedging Strategies for Energy Derivatives

29.1	Introduction	623
29.2	Literature.....	626
29.3	LRM Hedging Strategy	626
	Existence of the LRM Strategy • Determination of the LRM Strategy • Stochastic Volatility Models	
29.4	Application to Non-Traded Assets.....	634
	Strategy Derived from the delta Hedge • (Adjusted) LRM Hedging Strategy	
29.5	Numerical Results.....	637
	Set-Up • Different Strategies • Results	
29.6	Conclusion	642
	Acknowledgements	642
	References.....	642

Peter Leoni

Nele Vandaele

Michèle Vanmaele

In this article, we define a hedging strategy in a setting typical for the commodity market. Firstly, we prove the existence of the locally risk-minimizing (LRM) hedging strategy for payment streams in this setting. Next, a three-step procedure is described to determine the LRM hedging strategy. Then the procedure is illustrated for stochastic volatility models, as these models are a special case of the non-traded situation which frequently occurs in the commodity markets. Finally, we introduce the (adjusted) LRM hedging strategy in the non-traded setting and for this specific setting we numerically show the outperformance of this strategy compared with current market practice.

Keywords: Energy derivatives, Hedging strategies, Non-traded assets, Local risk minimization, Stochastic volatility models

JEL Classification: G13, G19

29.1 Introduction

Hedging under restrictions is a problem of great practical importance. It can become relevant in all financial markets but it is extremely important in energy markets where liquidity can be poor. Traditionally, a lot of academic attention has gone to setting up risk-neutral spot models that immediately model the short-to-delivery contracts, such as a day-ahead forward, under the assumption that hedging can be done perfectly. However, the market does not trade these short-to-delivery contracts very far ahead in time. In fact, most of the trading is focused on lower granular delivery periods such as calendar years or seasons. In this paper, we will focus on the dynamics of such forward contracts and establish how to hedge derivatives on smaller delivery periods, such as months or days by using a coarse-grained trading approach. Of course,

this approach introduces a basis risk as the claims are written on a different asset than the one that is being used to hedge with. But in energy markets, the extra risk is worth taking because the liquidity constraint would either not allow a direct hedging procedure to be executed, or the associated cost would be too high.

We will define an (adjusted) locally risk-minimizing (LRM) hedging strategy to a setting that is a good representation of everyday practice in an energy market. The obtained LRM hedging strategy is compared to a few common practices and shown to outperform significantly.

Since energy such as electricity or gas are non-storable commodities, the trading market has been organized around futures and forwards. These contracts provide an agreement between the two transacting parties to deliver the commodity over a fixed period of time rather than ensuring an instantaneous delivery. This means that the variety of delivery periods is enormous and although the correlation between them is not always strong, usually only a few contracts are liquid enough to execute trading strategies.

Since the underlying asset is a commodity that gets delivered physically in a certain volume, it is natural that the prices are denominated in currency per unit of volume per unit of time. Over the years, energy markets have attained a specific structure suitable for handling this flow nature. One of the unique features of the forward curve is its decomposition, or bucketing, into different granularities. Far ahead in the future, the only forward contracts traded are forwards for delivery of power over a complete calendar year. Once the calendar year approaches, these contracts gradually break down into quarterly contracts in a 'cascading' process. Closer to delivery, these quarterly contracts will break up into monthly, weekly and even daily forward contracts.

We discuss this setting in more detail for the electricity market and for the gas market separately.

- i. Electricity market: Besides the delivery period during which European power or electricity gets delivered, one often distinguishes three different products: base, peak and off-peak. Those are best explained by means of an example. A CAL-14 peak product is a contract that will deliver electricity during the entire calendar year 2014, but only during the peak hours of the day. That means that during the weekend no power will be delivered and during the weekdays, the delivery only takes place during the day (peak hours). A base contract ensures delivery of power during every single hour of the delivery period, without exception. The specific definition of which hour is a peak hour or off-peak hour, depends on the market in question.

As indicated, prices are denominated in currency per unit of volume per unit of time. For example, a CAL-14 base product for a volume of 10 MW (megawatt) could have a price of €50/MWh. This means that per delivered hour and per MW, the price is €50. There are 8760 h in the year 2014, so the total premium that is paid for the delivered power is $€50 \times 8760 \times 10 = €4\,380\,000$.

A peak contract for the same delivery period 2014 could have a price of €80/MWh, which is higher per active unit, but since the number of active hours is much less, the total premium is still lower compared to the base contract. Roughly speaking the number of peak hours is one-third of the amount of base hours. This means that the premium is $€80 \times 8760/3 \times 10 = €2\,336\,000$.

It is clear that there is a relationship between the peak, off-peak and base price in the market. If one buys electricity for delivery during peak hours and at the same time buys a contract that ensures delivery during the off-peak hours, it is obvious that the power is delivered without interruption and this is equivalent to a base contract.

If we denote the forward price of a peak contract by $F^{(p)}$ and the forward price of an off-peak contract as $F^{(o)}$, then the forward price of the base contract $F^{(b)}$ is given by

$$F^{(b)} = w^{(p)}F^{(p)} + w^{(o)}F^{(o)}$$

where the weights $w^{(p)}$ and $w^{(o)}$ depend on the number of peak and off-peak hours of the market. Since all prices are normalized to one unit of power and one hour, the typical weights are $w^{(p)} = 1/3$ and $w^{(o)} = 2/3$, where we should never forget that the actual cash flows will take into account the number of hours (see examples above).

In terms of liquidity, peak or off-peak contracts are not as liquid as base, especially far ahead in the future. This has its implications for the writer of an option on a peak (or off-peak) forward contract. Although it is one of the very basic assumptions in derivatives theory, energy market traders often find themselves in a situation where they sell options although the underlying contract is not liquidly traded. In the application of the theory, we will assume that a claim on peak power is transacted. It could be that at the time of this transaction, the value of this peak contract is known but that the spread between the bid price and the offer price is too big to efficiently hedge this position. For this reason, it is common practice to hedge with a base product rather than a peak product until closer to maturity of the option the strategy is swapped into a strategy in the peak product, when the liquidity has increased. We will call the peak product non-tradable for reasons of illiquidity and in [Section 29.5](#) we will study the effect of different hedging strategies.

- ii. Gas market: The gas market is highly seasonal with summer prices usually substantially cheaper than winter prices. Because of this, it is easy to understand that the most liquid products are forward contracts for the delivery of gas during the winter or during the summer. Gas by itself is more storable than power because the pressure differences in the network allow for the variations in demand during the day. Because of this, there is no peak and off-peak market for gas.

The interest in option contracts in the European gas market (e.g. UK or NBP market) is increasing, since these kind of contracts provide an interesting way of hedging the portfolios of big gas players in the market. However, for historical reasons, the most liquid option contract is a so-called seasonal option, that actually consists of a strip of 6 monthly options. So there are six underlying levels $F^{(1)}, F^{(2)}, \dots, F^{(6)}$ that are relevant for the pricing and hedging of such a contract. However, at the time when the options are written, not all of these monthly forward prices are known and the hedging has to be done by means of the season forward, which is in fact given by

$$F = \sum_{i=1}^6 w^{(i)} F^{(i)}. \quad (29.1)$$

In general, we will denote the underlyings as a vector namely $(F^{(1)}, \dots, F^{(d)})$ and the corresponding weights as $(w^{(1)}, \dots, w^{(d)})$.

As we will show in this article, it is even not possible to determine a delta hedge for these non-traded settings. Hence in practice adjusted delta hedges are used instead. We will propose an (adjusted) LRM hedging strategy and compare it to the adjusted delta hedges used in practice.

The next section contains an overview of articles dealing with non-traded assets. In [Section 29.3](#), we give a short introduction to the theory of (local) risk minimization. In [Section 29.3.1](#), we show how our setting fits into the very general setting found in literature and discuss the existence of a LRM hedging strategy for payment streams. Next, in [Section 29.3.2](#) we describe a three step procedure to compute this strategy for the setting valid for the commodity market. We refer the readers interested in the more technical details concerning the LRM hedging strategy to, e.g. Vandaele and Vanmaele (2008) and the references there in. We remark that stochastic volatility models are a special case of the non-traded setting, we work with. Therefore, in [Section 29.3.3](#) we determine the LRM hedging strategy for the stochastic volatility models studied in Poulsen *et al.* (2009), but following our three step approach. Another quadratic hedging strategy we could look at is the mean-variance hedging (MVH) strategy. This strategy has the advantage that it is self-financing and that it minimizes the total cost. Although theoretically speaking, the solution to this strategy is known, in practice it is still hard to find the exact solution. For more details concerning the problems arising in the non-traded asset case, we refer to Vandaele (2010). In the subclass of (affine) stochastic volatility models, explicit MVH strategies can be determined, see e.g. Kallsen and Vierthauer (2009) and Kallsen and Pauwels (2010).

Assuming specific dynamics for the forward price processes, the determination of the LRM and the adjusted LRM hedging strategies for the commodity setting described at the beginning of this section

is given in [Section 29.4](#). In [Section 29.5](#), we compare the different current market practices to the newly obtained results and show that the proposed adjusted LRM hedging strategies outperform, even in the simplest cases. Hence we show with numerical results that the framework of LRM hedging strategies is extremely relevant for managing the risks in contracts on non-traded assets. The flow of information from mature financial markets to energy markets has only just begun. This paper is to our knowledge the first to successfully apply advanced mathematical results to the everyday practice of hedging options in power and gas markets.

29.2 Literature

We found only few articles dealing with non-traded assets. The setting used in these articles differs from the one we work with, because they start from a different underlying problem. Furthermore, they all concentrate on the continuous setting, while we also allow discontinuous processes.

In literature, see e.g. Davis (2006), the term basis risk is often used for the non-hedgeable risk which remains and cannot be hedged away due to the fact that the asset on which the option is written is not available for hedging. Hedging in this case can only be done by using some closely related asset. Sometimes, the underlying asset is available for hedging but is too expensive due to transactions costs.

We mention here some important papers, the interested reader can also look at the references in those papers.

- Davis (2006) assumes that the underlying asset cannot be traded but is observable. Instead ‘a closely related’ asset, with a continuous price process, is traded. This closely related asset is assumed to follow a Brownian motion which is correlated with the underlying risky asset. The optimal hedging strategy is determined using exponential utility as a criterion. Numerical results for this case are derived in Monoyios (2004).
- Henderson (2002) and Henderson and Hobson (2002) work in the same setting. Furthermore, they also determine the power utility and give numerical results.
- Hobson (2005) gives an upper bound for the utility indifference price of a contingent claim on a non-traded asset in the same setting as described by Davis (2006).
- Ankirchner *et al.* (2010) also calculate the exponential utility-based indifference prices and corresponding hedges in a continuous setting. Their results are obtained in terms of solutions of forward–backward stochastic differential equations. Hence, the optimal hedging strategies are described in terms of the indifference price gradient and the correlation coefficients. Furthermore, the hedge can be seen as a generalization of the ‘delta-hedge’ in complete markets.
- Horst *et al.* (2010) concentrate on transferring nonfinancial risk, as for example depending on the temperature, to the capital markets. They give numerical results of equilibrium prices and optimal utilities in a continuous framework.
- Njoh (2007) determines quadratic hedging strategies for electricity spot prices with continuous price processes when another related asset is used for hedging.

We also wish to point out the growing interest in literature for numerical comparisons between the delta hedge and the quadratic hedging strategies, see Altmann *et al.* (2008), Denkl *et al.* (2009) and Brodén and Tankov (2011).

29.3 LRM Hedging Strategy

In this section, we concentrate on the existence and the determination of the LRM hedging strategy for payment streams in terms of Radon–Nikodym derivatives by a three-step procedure. Then, we illustrate our procedure by applying it to stochastic volatility models, as these models are a special case of the non-traded situation which frequently occur in the commodity market.

In Section 29.3.1, we first state the results ensuring the existence of the Föllmer–Schweizer (FS) decomposition based on Choulli *et al.* (1998) and the existence of the LRM hedging strategy in a multidimensional setting and allowing for payment streams using Schweizer (2008). Assuming all the conditions for the existence of the LRM hedging strategy, we determine in Section 29.3.2 the optimal number of risky assets based on Choulli *et al.* (2010), where an explicit form for the FS decomposition is given in terms of the predictable characteristics. In fact the setting used in Choulli *et al.* (2010) is much more general than is needed in the present paper. The goal in that article is the determination of the FS decomposition in a framework as general as possible, while the link with the LRM hedging strategy was not described explicitly. The LRM hedging strategy in continuous time only makes sense if the finite variation part is continuous and hence we make this restriction in our setting. Due to this assumption, we can obtain very quickly a shorthand notation for the optimal number of risky assets, which will certainly help the reader to understand the objective of this strategy and the FS decomposition. We remark that the way we follow here is more linked with some of the intermediate results given in Černý and Kallsen (2007) on their way to determining the MVH strategy.

LRM hedging strategies control less of the total risk than mean-variance ones, but the advantage is that solutions for the hedging strategies are much simpler. In the martingale setting, there is no difference between the optimal number of risky assets given by the risk-minimizing (RM) and the MVH strategy. Both solutions are found through the Galtchouk–Kunita–Watanabe (GKW) decomposition.

For more details concerning quadratic strategies, we refer to e.g. Pham (2000), Schweizer (2001), Vandaele and Vanmaele (2008) and the references therein. We also remark that the results of this section can be seen as an extension of the work by Colwell and Elliott (1993). Colwell and Elliott looked at the case when the contingent claim depends solely on the risky assets in which one can invest, but at the time the article was written the theory of local risk minimization was less developed. We felt that it would be possible to describe an easier, more straightforward approach to tackle not only the setting described by Colwell and Elliott (1993), but an even wider range of problems, including the problem of non-traded assets and payment streams. Choulli *et al.* (2010) extended Colwell and Elliott (1993) in the sense that they described a more direct way to describe the FS decomposition and the dynamics used to describe the process of the risky asset are far more general.

To further illustrate the usefulness of the method described here, we also refer to Poulsen *et al.* (2009). They determine the LRM hedging strategy for a general class of stochastic volatility models driven by continuous processes using the three-step procedure introduced by El Karoui *et al.* (1997). Following that approach, the market is first completed, then the hedging strategy is computed in the completed market and finally this strategy is projected on the original market. In Section 29.3.3, we obtain the same solution using our three step approach described here, which does not require the continuity of the process.

29.3.1 Existence of the LRM Strategy

We work on the probability space $(\Omega, \mathcal{F}, P)$. The filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ satisfies all the usual conditions and $T \in [0, +\infty)$ is the fixed time horizon.

When the underlyings of the contingent claim that one wants to hedge are semimartingales we can no longer apply the RM hedging strategy, but we need instead to apply the LRM hedging strategy described for the first time by Schweizer in a series of articles. The goal of a RM hedging strategy is to minimize the variance of the cost process C , while the portfolio should equal the contingent claim at payment date. Hence, we minimize at any time t

$$E[(C_T - C_t)^2 | \mathcal{F}_t],$$

with $C_t = V_t - \int_0^t \xi_u dS_u$ where V denotes the value process of the portfolio and $\int_0^g \xi_u dS_u$ is the gain process of trading in the underlying S . The exact criterion, which is minimized in the case of semimartingales, becomes a bit more involved. Therefore, we refer the interested reader to, e.g. Schweizer (1991). The RM hedging strategy is easily derived from the GKW decomposition:

Definition 29.3.1 Galtchouk–Kunita–Watanabe (GKW) decomposition: If S is a locally square integrable local martingale under the measure Q , then

$$H = H^{(0)} + \int_0^T \xi_u^H dS_u + L_T^H \quad Q\text{-a.s.}$$

is the GKW-decomposition of the square-integrable contingent claim H if $\xi^H \in L(S)$, if $\xi^H \cdot S$ and L^H are local Q -martingales and if L^H is orthogonal to S , with $L_0^H = 0$.

Here we used the shorthand notation for stochastic integrals, namely, $\xi \cdot S$ stands for $\int_0^T \xi_u dS_u$. The optimal amount invested in the risky asset is given by ξ^H in Definition 29.3.1.

For the LRM hedging strategy, we define first the space Θ_S which we need in the definition of the FS decomposition.

Definition 29.3.2 The space Θ_S consists of all $\mathbb{R}^d$ -valued predictable processes θ such that the stochastic integral process $\theta \cdot S$ is well-defined and belongs to the space $S^2(P)$ of semimartingales. This means that

$$E \left[\int_0^T \theta_s' d\langle M \rangle_s \theta_s + \left(\int_0^T \theta_s' dA_s \right)^2 \right] < \infty,$$

where $'$ denotes the transpose, $\langle \cdot \rangle$ is the predictable quadratic variation, M is the martingale part and A the finite variation part of S .

The LRM hedging strategy is found through the FS decomposition:

Definition 29.3.3 Föllmer–Schweizer (FS) decomposition: If S is a semimartingale under the measure P , then

$$H = H^{(0)} + \int_0^T \xi_u^H dS_u + L_T^H \quad P\text{-a.s.}$$

is the Föllmer–Schweizer decomposition of the contingent claim H if the S -integrable process $\xi^H \in \Theta_S$ and if L^H is a square-integrable P -martingale orthogonal to the martingale part M of S , with $L_0^H = 0$.

For continuous semimartingales the FS decomposition can be deduced from the GKW-decomposition under the *minimal martingale measure* (MMM), which is the unique measure such that every local martingale L under P , which is also orthogonal to the martingale part M of S remains a martingale under the MMM. This approach is no longer valid for discontinuous semimartingales.

In Choulli *et al.* (2010), the relationship between the hedgeable part of the GKW decomposition under the MMM and the one of the FS decomposition under the original measure is given and the FS decomposition is determined using the predictable characteristics. In the next section, we will show how we can determine the decomposition following a three step procedure in terms of Radon–Nikodym derivatives which is implicitly mentioned in Choulli *et al.* (2010). The aim here is to show that our setting fits in the very general setting of Choulli *et al.* (2010) and to derive the existence conditions for the LRM strategy from it. Since we concentrate here on the determination of the LRM strategy this less general setting, which is typical for the commodity market, is not really restrictive. Furthermore, we derive the hedging strategy for payment streams based on Schweizer (2008) where it is proved that even in the multidimensional case and for payment streams the LRM hedging strategy is determined by the FS decomposition. In that paper also some conditions on the hedgeable risky asset S as given in Schweizer (1991) are relaxed.

Definition 29.3.4 A $\mathbb{R}^d$ -valued semimartingale $S = (S_t)_{0 \leq t \leq T}$ satisfies the *structure condition* (SC) if S is special with canonical decomposition

$$S = S_0 + M + A = S_0 + M + \int d\langle M \rangle \lambda,$$

where M is a square integrable local martingale with $M_0 = 0$ and λ is a $\mathbb{R}^d$ -valued, predictable process that is locally square integrable w.r.t. M , so that the *mean-variance trade-off process* $K := \int \lambda' dA = \int \lambda' \langle M \rangle \lambda$ satisfies $K_T < \infty$ P -a.s.

We remark that the SC is related to the absence of arbitrage condition.

Definition 29.3.5 An L^2 -strategy is a pair $\phi = (\xi, \tilde{\eta})$, where $\xi \in \Theta_S$ and $\tilde{\eta}$ is a real-valued adapted process such that the value process $V(\phi) := \xi S + \tilde{\eta}$ is right-continuous and square-integrable.

Theorem 29.3.6: Schweizer (2008)

Suppose the $\mathbb{R}^d$ -valued semimartingale S satisfies the SC and the mean-variance trade-off process $K = \int \lambda' d\langle M \rangle \lambda$ (or, equivalently, A) is continuous. Then a payment stream H admits a LRM L^2 -strategy ϕ if and only if H_T admits a FS decomposition. In that case $\phi = (\xi, \tilde{\eta})$ is given by

$$\xi = \theta^{H_T}, \quad \tilde{\eta} = V^{H_T} - \theta^{H_T} S$$

with

$$V_t^{H_T} := H_t^{(0)} + \int_0^t \theta_s^{H_T} dS_s + L_t^{H_T} - H_t, \quad 0 \leq t \leq T,$$

where the processes with superscript H_T are coming from the FS decomposition of H_T .

It is important to notice that this theorem guarantees that if you find a LRM hedging strategy then this strategy makes sense and will really minimize the risk locally, but this does not guarantee the existence of this strategy.

Here to we have to combine this with the existence of the FS decomposition as proved in Choulli *et al.* (1998) under certain conditions for a very general class of so-called ε -martingales. In our setting, the Girsanov density process Z describing the change of measure from P to the MMM $\tilde{Q}$ is given by $\varepsilon(N)$ the Doléans-Dade exponential of $N = -\lambda \cdot M$ which is a special case of Choulli *et al.* (1998) where most of the proofs are given for an arbitrary N . From Proposition 3.7 in Choulli *et al.* (1998), we know that if $\langle N \rangle_T \in L^\infty$ then the conditions for the existence of a FS decomposition are satisfied. If we assume moreover that this process Z is a strictly positive square-integrable P -martingale then the class of ε -martingales simplifies to $\tilde{Q}$ -martingales. Hence, we can conclude with the following theorem which combines the main result from Schweizer (1991) and Choulli *et al.* (1998):

Theorem 29.3.7

Suppose S is an $\mathbb{R}^d$ -valued special semimartingale satisfying the SC and whose finite variation part A is continuous. If also $\langle \lambda \cdot M \rangle_T \in L^\infty$ then the LRM hedging strategy for a payment stream H and the FS decomposition of H_T exist, and the L^2 -strategy in which we should invest can be found from the FS decomposition.

An important underlying condition, which is now hidden beneath other conditions, is the closedness of the space $G_T(\Theta_S) := \{(\theta \cdot S)_T : \theta \in \Theta_S\}$. Monat and Stricker (1995) proved that if $\langle \lambda \cdot M \rangle_T \in L^\infty$ then the space $G_T(\Theta_S)$ is closed in L^2 . Choulli *et al.* (1998) generalized this result to the class of ε -martingales.

29.3.2 Determination of the LRM strategy

Based on the previous section the following conditions are assumed to be satisfied:

- S is a special semimartingale satisfying the SC;
- the finite variation part A is continuous;
- $\langle \lambda \cdot M \rangle_T \in L^\infty$, with M the P -martingale part of S ;
- the density process $Z = \varepsilon(-\lambda \cdot M)$ is a strictly positive square-integrable P -martingale.

The procedure to determine the LRM hedging strategy is based on the proof of Theorem 5.5 of Choulli *et al.* (1998) but is much simpler because of the assumptions that we made. We denote by H the payment process for which we want to determine the hedging strategy.

Step 1 Define the $\tilde{Q}$ -martingale Y with $Y_T = H_T$ by

$$Y_t = \frac{E[H_T Z_T | \mathcal{F}_t]}{Z_t} = E^{\tilde{Q}}[H_T | \mathcal{F}_t] \text{ on } \{t \in [0, T]\}. \quad (29.3)$$

The value of the hedging portfolio V^H at time t is then the difference of Y and the payment process H at t , i.e. $V_t^H = Y_t - H_t$.

Comment: The existence and the form of this Y is given by Proposition 3.12 (iii) of Choulli *et al.* (1998) but in a simplified form thanks to the assumption that Z never becomes zero. In the case of a contingent claim with pay-off H at T the payments at times $t \in [0, T)$ are zero and the value of the hedging portfolio V^H at time t equals Y_t .

Step 2 Extract the P -martingale part I of the $\tilde{Q}$ -martingale Y which equals

$$I = Y - Y^{(0)} + \langle I, N \rangle^P$$

where $\langle \cdot, \cdot \rangle^P$ stands for the quadratic covariation of two P -martingales.

Comment: Note that we made here the dependence of the quadratic (co)variation process on the measure explicit by using a superscript. Furthermore we know that $-\langle I, N \rangle^P$ is the finite variation part of the $\tilde{Q}$ -martingale Y under P , see also Proposition 4.2 in Choulli *et al.* (2010).

From the FS decomposition of H_T we may infer $Y = Y^{(0)} + \xi \cdot S + L$ and hence $I = \xi \cdot M + L$ which is the GKW decomposition of the P -martingale I .

In the setting described here the finite variation part of Y is thus also continuous because $-\langle I, N \rangle^P = \langle \xi \cdot M + L, N \rangle^P = \lambda \xi \cdot \langle M, M \rangle^P = \xi \cdot A$, see Definition 29.3.4.

The problem of finding the FS decomposition of H_T is now reduced to the determination of the GKW decomposition of the P -martingale I .

Step 3 Compute the optimal number ξ as the Radon–Nikodym derivative:

$$\xi = (d\langle M, M \rangle^P)^{\text{inv}} d\langle I, M \rangle^P.$$

The strategy at time t is then given by $(\xi_t, \tilde{\eta}_t)$ with $\tilde{\eta}_t = Y_t - \xi_t S_t - H_t$.

Comment: The superscript inv denotes the inverse of the matrix. Here we mean the pseudoinverse of Moore–Penrose, see Albert (1972), whose existence is guaranteed.

The expression (29.3) follows by taking the angle bracket process with respect to M under P of both sides of $I = \xi \cdot M + L$ and solving for ξ :

$$\langle I, M \rangle^P = \xi \cdot \langle M, M \rangle^P.$$

To understand the difference between the LRM and the RM hedging strategy under the MMM $\tilde{Q}$, we also calculate the RM strategy:

$$\tilde{\xi}_t = (d\langle S, S \rangle_t^{\tilde{Q}})^{\text{inv}} d\langle Y, S \rangle_t^{\tilde{Q}} \quad (29.4)$$

and $\tilde{\eta}_t = Y_t - \tilde{\xi}_t S_t - H_t$. Note that when S is not a $\tilde{Q}$ -martingale, $\langle S, S \rangle^{\tilde{Q}}$ has to be understood as the quadratic variation of the $\tilde{Q}$ -martingale part of S . Similarly, $\langle Y, S \rangle^{\tilde{Q}}$ is the quadratic covariation of the $\tilde{Q}$ -martingale part of S and the $\tilde{Q}$ -martingale Y .

Hence, one can easily determine the relationship between the GKW and the FS decomposition in the following way: starting from Definition 29.3.1 applied to Y , we find

$$\langle Y, S \rangle^{\tilde{Q}} = \tilde{\xi} \cdot \langle S, S \rangle^{\tilde{Q}} + \langle L^{\text{GKW}}, S \rangle^{\tilde{Q}} = \tilde{\xi} \cdot \langle S, S \rangle^{\tilde{Q}}. \quad (29.5)$$

On the other hand, starting from Definition 29.3.3 we have that

$$\langle Y, S \rangle^{\tilde{Q}} = \xi \cdot \langle S, S \rangle^{\tilde{Q}} + \langle L^{\text{FS}}, S \rangle^{\tilde{Q}}.$$

Solving for ξ and using (29.5), we obtain

$$\begin{aligned} \xi &= (d\langle S, S \rangle^{\tilde{Q}})^{\text{inv}} d\langle Y, S \rangle^{\tilde{Q}} - (d\langle S, S \rangle^{\tilde{Q}})^{\text{inv}} d\langle L^{\text{FS}}, S \rangle^{\tilde{Q}} \\ &= \tilde{\xi} - \Phi. \end{aligned}$$

In this form, we obviously see that there is no equivalence between the GKW decomposition and the FS decomposition when the orthogonality between L^{FS} and M under P , is not preserved under $\tilde{Q}$. For an explicit counterexample where $\Phi = (d\langle S, S \rangle^{\tilde{Q}})^{\text{inv}} d\langle L^{\text{FS}}, S \rangle^{\tilde{Q}} \neq 0$, we refer to Choulli *et al.* (2010). Furthermore, $\Phi = 0$ in the continuous case, because for continuous processes the orthogonality is always preserved.

29.3.3 Stochastic Volatility Models

If the liquidity of the market increases, the liquidity in the derivatives follows as well. In an established option's market, options with various strikes all become readily traded and one can observe smile and/or skew behaviour in the implied volatilities. It is well known by now that this smile behaviour of long-dated options can be well explained by the introduction of another stochastic factor in the volatility or variance process. In particular, options on oil-related markets such as Brent or WTI exhibit this kind of strike-dependent volatility and therefore in this context it is natural to introduce a stochastic volatility model. The MVH strategy of claims on risky assets with a stochastic volatility is often investigated in literature. As references we mention Biagini *et al.* (2000), Černý and Kallsen (2008), Chan *et al.* (2009), Kallsen and Vierthauer (2009), Kallsen *et al.* (2010) and Kallsen and Pauwels (2010, 2011). The references to the LRM hedging strategy for stochastic volatility models are more rare: Frey and Runggaldier (1999) and Poulsen *et al.* (2009).

In Bertsimas *et al.* (2001), the ϵ -arbitrage strategy is determined for stochastic volatility models.

We illustrate here the technique for the determination of the LRM strategy as described in [Section 29.3.2](#) on the class of stochastic volatility models used in Poulsen *et al.* (2009). They apply a three-step procedure introduced by El Karoui *et al.* (1997): the market is completed, the hedging strategy is calculated in this completed market and then projected on the original market. In fact our procedure is similar, in that the determination of the MMM is part of completing the market. However, our approach is more general because the approach of El Karoui *et al.* (1997) is limited to Brownian motions. Furthermore, we also know that in the continuous setting the FS decomposition can easily be deduced from the GKW decomposition. This will be shown explicitly on the example applied in Poulsen *et al.* (2009).

The hedge is determined for European claims $H(S_T^*)$ when the underlying undiscounted risky asset S^* follows a stochastic volatility model of the following form:

$$\begin{aligned}\frac{dS^*(t)}{S^*(t)} &= \mu dt + S^*(t)^\gamma f(V(t)) \left[\sqrt{1-\rho^2} dW^1(t) + \rho dW^2(t) \right] \\ \frac{dV(t)}{V(t)} &= \beta(V(t))dt + g(V(t))dW^2(t),\end{aligned}$$

with independent standard Brownian motions W^1 and W^2 . We refer to Poulsen *et al.* (2009) for an overview of the models contained in this class and for more details concerning the functions/parameters β , γ , μ , ρ , g and f .

With S we denote the discounted dynamics:

$$\frac{dS(t)}{S(t)} = (\mu - r)dt + S^*(t)^\gamma f(V(t)) \left[\sqrt{1-\rho^2} dW^1(t) + \rho dW^2(t) \right],$$

where r denotes the risk-free interest rate. The notation M is used for the martingale part of the risky asset S . We are in a Markovian market model and we can easily determine the MMM

$$\tilde{Q}: Z_t = E\left[\frac{d\tilde{Q}}{dP} \middle| \mathcal{F}_t\right] = \varepsilon(-\lambda \cdot M)_t \text{ with}$$

$$\begin{aligned}dM(t) &= S(t)S^*(t)^\gamma f(V(t)) \left[\sqrt{1-\rho^2} dW^1(t) + \rho dW^2(t) \right] \\ d\langle M, M \rangle_t^P &= S(t)^2 S^*(t)^{2\gamma} f^2(V(t))dt \\ \lambda_t &= \frac{\mu - r}{S(t)S^*(t)^{2\gamma} f^2(V(t))}.\end{aligned}$$

Therefore $Z_t = \mathcal{E}\left[-\frac{\mu - r}{S^*(t)^\gamma f(V(t))}(\sqrt{1-\rho^2} dW^1(t) + \rho dW^2(t))\right]$. We start by deriving the hedging strategy for the claim $H(S_T^*)$ on the basis of the GKW decomposition.

Step 1 We define the $\tilde{Q}$ -martingale Γ according to (2). By the Markovian property, we know that Γ at time t only depends on $S^*(t)$ and $V(t)$:

$$\Gamma(t, S^*(t), V(t)) = e^{-r(T-t)} E^{\tilde{Q}}[H(S^*(T)) | S^*(t), V(t)]. \quad (29.6)$$

Applying Itô's formula to this $\tilde{Q}$ -martingale gives

$$d\Gamma(t, S^*(t), V(t)) = \Gamma_{S^*(t)} d(S^*)^{m,Q}(t) + \Gamma_{V(t)} V(t) g(V(t)) dW^{2,Q}(t),$$

with $W^{2,\tilde{Q}}$ the $\tilde{Q}$ -Brownian motion originating from W^2 and with $(S^*)^{m,\tilde{Q}}$ the $\tilde{Q}$ -martingale part of S^* with the following dynamics

$$d(S^*)^{m,Q}(t) = S^*(t)^{\gamma+1} f(V(t)) \left[\sqrt{1-\rho^2} dW^{1,Q}(t) + \rho dW^{2,Q}(t) \right],$$

with $W^{1,\tilde{Q}}$ the $\tilde{Q}$ -Brownian motion originating from W^1 .

Step 2 $I = \Gamma$ since we work under the MMM $\tilde{Q}$.

Step 3 The angle bracket process of S under $\tilde{Q}$ is given by

$$d\langle S, S \rangle_t^{\tilde{Q}} = S(t)^2 S^*(t)^{2\gamma} f(V(t))^2 dt.$$

The number of risky assets invested in S originating from the GKW decomposition is in view of 29.5 given by

$$\begin{aligned} \xi_t^{\text{GKW}} &= \frac{d\langle \Gamma, S \rangle_t^Q}{d\langle S, S \rangle_t^Q} = \frac{\Gamma_{S^*(t)} d\langle (S^*)^{m,Q}, S \rangle_t^Q + \Gamma_{V(t)} S^*(t)^\gamma S(t) g(V(t)) f(V(t)) V(t) \rho dt}{d\langle S, S \rangle_t^Q} \\ &= \left(\Gamma_{S^*(t)} + \rho \frac{V(t) g(V(t))}{S^*(t)^{\gamma+1} f(V(t))} \Gamma_{V(t)} \right) \frac{S^*(t)}{S(t)}. \end{aligned} \quad (29.7)$$

Hence, the amount we have to invest in S^* exactly equals the one given by Poulsen *et al.* (2009):

$$\Gamma_{S^*} + \rho \frac{V(t) g(V(t))}{S^*(t)^{\gamma+1} f(V(t))} \Gamma_{V(t)},$$

where the first term is the amount we have to invest according to the delta hedge and the second term can be seen as a correction term. Thus the hedging strategy in this case is $(\xi^{\text{GKW}}, \tilde{\eta})$ with $\tilde{\eta} = \Gamma_t - \xi_t^{\text{GKW}} S_t$.

Next, we determine the LRM strategy for the claim under the original measure P .

Step 1 This is the same as Step 1 above.

Step 2 The dynamics of the P -martingale part I of Γ (29.6) are

$$dI(t) = \Gamma_{S^*(t)} d\langle (S^*)^{m,P}, M \rangle_t^P + \Gamma_{V(t)} V(t) g(V(t)) dW^2(t),$$

where

$$d\langle (S^*)^{m,P}, M \rangle_t^P = \frac{S^*(t)}{S(t)} dM(t).$$

Step 3 Taking the P -angle bracket of M with I gives

$$d\langle I, M \rangle_t^P = \Gamma_{S^*(t)} d\langle (S^*)^{m,P}, M \rangle_t^P + \Gamma_{V(t)} V(t) g(V(t)) d\langle W^2, M \rangle_t^P,$$

with

$$d\langle (S^*)^{m,P}, M \rangle_t^P = \frac{S^*(t)}{S(t)} d\langle M, M \rangle_t^P$$

and

$$d\langle W^2, M \rangle_t^P = S^*(t)^\gamma S(t) f(V(t)) \rho dt.$$

Therefore the units of stock for the LRM hedging strategy equal

$$\xi_t^{\text{FS}} = \frac{d\langle I, M \rangle_t^P}{d\langle M, M \rangle_t^P} = \left(\Gamma_{S^*(t)} + \rho \frac{V(t) g(V(t))}{S^*(t)^{\gamma+1} f(V(t))} \Gamma_{V(t)} \right) \frac{S^*(t)}{S(t)},$$

which is exactly the number found from the determination of the GKW decomposition given in (29.7). The LRM strategy is $(\xi^{\text{FS}}, \tilde{\eta})$ with $\tilde{\eta} = \Gamma_t - \xi_t^{\text{FS}} S_t$.

This demonstrates again the equality between the GKW decomposition under the MMM and the FS decomposition under the original measure in the continuous case.

We remark that we can easily extend this model in several ways. We can add for example jumps, a stochastic interest rate or use a vector to describe the stochastic volatility. We will not do this explicitly for this setting, because all these stochastic volatility models can be seen as a special case of non-traded assets. Namely choose for $F^{(1)}$ the risky asset and for $(F^{(2)}, \dots, F^{(d)})$ the (vector of) stochastic volatilities, furthermore as weights we take the vector $(1, 0, \dots, 0)$. For an example in which the LRM hedging strategy for the Bates model is calculated, we refer to Hubalek and Sgarra (2007).

29.4 Application to Non-Traded Assets

The *fixed* combination in which we can invest will be denoted by $F := \sum_{i=1}^d w^{(i)} F^{(i)}$, with $F^{(i)}$ the forward prices. We remark that the filtration $\mathbb{F}$ contains the information of the non-traded assets with prices $F^{(i)}$, $i = 1, \dots, d$, as well as the information of the traded asset with price F , in the sense that the filtration is generated by the underlying driving processes. By Γ , we denote the price of the claim written on the non-traded assets under a martingale measure which we do not specify here yet.

29.4.1 Strategy Derived from the Delta Hedge

In practice, non-traded assets are often hedged using a strategy based on the delta hedge. We will use the intuitively obtained hedging strategies to compare them with the LRM hedging strategy.

In this section, we restrict to the two-dimensional case for illustrations but the discussion is easily extended to higher dimensions.

In the standard two-dimensional case, we trade $\xi^{(i)}$ forwards with price $F^{(i)}$, $i = 1, 2$, such that the risk originating from the rate of change of the claim price with respect to these forward prices equals

$$\left[\frac{\partial \Gamma}{\partial F^{(1)}} - \xi^{(1)} \right] dF^{(1)} + \left[\frac{\partial \Gamma}{\partial F^{(2)}} - \xi^{(2)} \right] dF^{(2)}. \quad (29.8)$$

This risk can be completely eliminated by choosing $\xi^{(i)}$ equal to $\frac{\partial \Gamma}{\partial F^{(i)}}$, $i = 1, 2$.

For non-traded assets, we can only invest in ξ assets with price F . It is impossible to eliminate the risk exposure completely because the following equations should be satisfied by ξ

$$\xi^{(1)} = w^{(1)} \xi \quad \text{and} \quad \xi^{(2)} = w^{(2)} \xi$$

and so we must search for the most optimal ξ . We give some intuitively-based solutions for ξ :

- Total volume-neutral strategy:

The number ξ to invest in F equals

$$\xi = \max \left(\frac{\partial \Gamma}{\partial F^{(1)}}, \frac{\partial \Gamma}{\partial F^{(2)}} \right).$$

- Volume-neutral strategy:

The power market has two natural units of volume since the commodity is delivered in a certain magnitude over a period of time. The magnitude is expressed in MW and the time in hours. So instead of focusing on the MW position, one could also focus on the total volume, taking into account the length of the delivery period:

$$\xi = w^{(1)} \frac{\partial \Gamma}{\partial F^{(1)}} + w^{(2)} \frac{\partial \Gamma}{\partial F^{(2)}}.$$

- Price-adjusted strategy:

$$\xi = \frac{w^{(1)} \frac{\partial \Gamma}{\partial F^{(1)}} F^{(1)} + w^{(2)} \frac{\partial \Gamma}{\partial F^{(2)}} F^{(2)}}{w^{(1)} F^{(1)} + w^{(2)} F^{(2)}}.$$

- Delta hedging with minimal risk exposure:

Restricting (29.8) to the setting of non-traded assets and calculating the differential of the portfolio consisting of the claim and ξ forwards with price F , we find the following:

$$\left[\frac{\partial \Gamma}{\partial F^{(1)}} - \xi w^{(1)} \right] dF^{(1)} + \left[\frac{\partial \Gamma}{\partial F^{(2)}} - \xi w^{(2)} \right] dF^{(2)}.$$

Here we have left out the dt -part because this part is not risky.

The variance of this remaining risk is in vector notation:

$$\text{var}(\xi) = \left[\frac{\partial \Gamma}{\partial F} - \xi \mathbf{w} \right] d\langle F, F \rangle^p \left[\frac{\partial \Gamma}{\partial F} - \xi \mathbf{w} \right],$$

where $\frac{\partial \Gamma}{\partial F}$ is the gradient of Γ , $\mathbf{w}$ is the vector containing the weights and $F := (F^{(1)}, F^{(2)})$. We minimize this variance to obtain the optimal ξ :

$$\frac{d \text{var}(\xi)}{d\xi} = -\mathbf{w}' d\langle F, F \rangle^p \left[\frac{\partial \Gamma}{\partial F} - \xi \mathbf{w} \right] - \left[\frac{\partial \Gamma}{\partial F} - \xi \mathbf{w} \right]' d\langle F, F \rangle^p \mathbf{w} = 0.$$

Solving this equation for this ξ , gives

$$\xi = \frac{\frac{\partial \Gamma}{\partial F} d\langle F, F \rangle^p \mathbf{w}}{\mathbf{w}' d\langle F, F \rangle^p \mathbf{w}}.$$

This is exactly the result we will obtain when we apply the LRM hedging theory when Γ is the price of the claim under the MMM linked with F . So we achieved here an intuitive explanation for the rather complicated theory of local risk minimization. We remark that we cannot follow blindly this intuitive approach, because in this way it is for example not clear why we have to use the MMM to price the claim.

We note that Poulsen *et al.* (2009) independently made an analogous conclusion.

These intuitively-based solutions will be used for comparison with the solutions to the local risk minimization.

29.4.2 (Adjusted) LRM Hedging Strategy

The theoretical determination of the LRM hedging strategy in this context is rather straightforward once one understands properly the existence conditions. A major concept for LRM hedging strategies is the MMM. Concentrating on our setting, it is important to understand that one must consider martingale measures with respect to the underlying asset which is used for hedging and NOT with respect to the assets on which the claim depends. More concretely, we need to determine the MMM of F and not of $F^{(1)}, \dots, F^{(d)}$. This illustrates the incompleteness of the market we work in, even in the case with two driving Brownian motions (there is only one martingale measure to be determined for the vector $(F^{(1)}, F^{(2)})$ from infinitely many for F). Due to the results described in [Section 29.3.2](#), the determination of the

LRM hedging strategy for non-traded assets is reduced to a straightforward application of the formulas given there.

We distinguish two different cases: the continuous case and the discontinuous case.

- In the continuous case, the LRM hedging strategy equals that of the RM hedging strategy under the MMM $\tilde{Q}$, see (29.4). Hence, the optimal position in the hedging asset F is given by

$$\xi = \frac{d\langle \Gamma, F \rangle^{\tilde{Q}}}{d\langle F, F \rangle^{\tilde{Q}}} = \frac{\frac{\partial \Gamma'}{\partial F} d\langle F, F \rangle^{\tilde{Q}} \mathbf{w}}{\mathbf{w}' d\langle F, F \rangle^{\tilde{Q}} \mathbf{w}}.$$

To see this, apply Itô's formula to Γ relying on its continuity and martingale property.

- In the discontinuous case, we apply formula (29.3) and hence the optimal position in the hedging asset F is given by

$$\xi = \frac{d\langle I, M \rangle^P}{d\langle M, M \rangle^P} = \frac{d\langle I, F^m \rangle^P \mathbf{w}}{\mathbf{w}' d\langle F^m, F^m \rangle^P \mathbf{w}}, \quad (29.9)$$

with I the P -martingale part of Γ and where M and F^m stand for the P -martingale part of F , respectively F .

To find the explicit position in risky assets, we need to fill in the dynamics of the processes in the brackets. Due to the typical setting it is not possible to express the optimal amount in terms of the cumulant function by using Fourier transformation as Hubalek *et al.* (2006) did, see Vandaele (2010) for more details. Therefore, we proceed in the following way:

If the forwards have the following dynamics under the original measure P

$$dF_t^{(i)} = F_t^{(i)} (\alpha(i)dt + \sigma^{(i)}dW_t^{(i)}), \quad i = 1, 2$$

then by applying Itô's formula to the expected price Γ of the claim, taken under the MMM of $F = w^{(1)}F^{(1)} + w^{(2)}F^{(2)}$, we obtain the dynamics of the P -martingale part I of Γ :

$$dI_t = \Gamma_{f^{(1)}} \sigma_t^{(1)} F_t^{(1)} dW_t^{(1)} + \Gamma_{f^{(2)}} \sigma_t^{(2)} F_t^{(2)} dW_t^{(2)},$$

where $\Gamma_{f^{(i)}}$ stands for $\frac{\partial \Gamma}{\partial F^{(i)}}$, $i = 1, 2$. Inserting these dynamics in (29.3) leads to the following value for ξ :

$$\xi = \frac{w^{(1)} \sigma^{(1)} F^{(1)} (\Gamma_{f^{(1)}} \sigma^{(1)} F^{(1)} + \Gamma_{f^{(2)}} \rho \sigma^{(2)} F^{(2)}) + w^{(2)} \sigma^{(2)} F^{(2)} (\Gamma_{f^{(1)}} \rho \sigma^{(1)} F^{(1)} + \Gamma_{f^{(2)}} \sigma^{(2)} F^{(2)})}{w^{(1)} (\sigma^{(1)} F^{(1)})^2 + 2w^{(1)} w^{(2)} \rho \sigma^{(1)} F^{(1)} \sigma^{(2)} F^{(2)} + w^{(2)} (\sigma^{(2)} F^{(2)})^2},$$

where ρ represents the correlation between the two Brownian motions $W^{(1)}$ and $W^{(2)}$.

We apply this exact formula when $F^{(1)}$ and $F^{(2)}$, and hence, also F , are martingales under the original measure. If F is not a martingale under the original measure then claim prices should be calculated under the MMM linked with F . To this end, a partial differential equation (PDE) has to be solved. However, the determination of this claim price by solving a PDE at every trading date is really time-consuming. Therefore, we propose an adjustment of the amount ξ , which can be calculated fast without solving a PDE. Namely, we use here as martingale measure not the MMM, but the unique martingale measure which ensures that both $F^{(1)}$ and $F^{(2)}$ become martingales. Hence, the optimal value for ξ in the *adjusted* LRM hedging strategy is given by

$$\xi = \frac{w^{(1)} \sigma^{(1)} F^{(1)} (\Delta^{(1)} \sigma^{(1)} F^{(1)} + \Delta^{(2)} \rho \sigma^{(2)} F^{(2)}) + w^{(2)} \sigma^{(2)} F^{(2)} (\Delta^{(1)} \rho \sigma^{(1)} F^{(1)} + \Delta^{(2)} \sigma^{(2)} F^{(2)})}{w^{(1)} (\sigma^{(1)} F^{(1)})^2 + 2w^{(1)} w^{(2)} \rho \sigma^{(1)} F^{(1)} \sigma^{(2)} F^{(2)} + w^{(2)} (\sigma^{(2)} F^{(2)})^2}, \quad (29.10)$$

where the amounts $\Delta^{(i)} := \frac{\partial \Gamma}{\partial F^{(i)}}$ are calculated using the martingale measure for $F^{(1)}$ and $F^{(2)}$ separately.

In fact they should be calculated under the MMM related to F in the semimartingale case. Still we find that this makes sense in our setting, because after some time $F^{(1)}$ and $F^{(2)}$ will become both liquid on the market and then the martingale measure used here is the correct one. Furthermore as we will discuss later on, finding the correct drift is almost impossible, hence finding the correct minimal martingale measure is equally difficult.

In the presence of jumps in the dynamics of $F^{(i)}$, $i = 1, 2$, the claim price solves a partial integro differential equation (PIDE). So it becomes even more involved. Hence, we will also use formula (29.10) in the discontinuous case, but then ρ no longer stands for the correlation between the two Brownian motions. Instead, ρ is the parameter used to express the correlation between the two forwards with prices $F^{(1)}$ and $F^{(2)}$. Furthermore, the martingale measure for $F^{(1)}$ and $F^{(2)}$ is no longer unique because of the incompleteness and we choose to work under the mean-correcting measure (see e.g. Schoutens (2003)) for $F^{(1)}$ and $F^{(2)}$ separately to calculate the $\Delta^{(i)} := \frac{\partial \Gamma}{\partial F^{(i)}}$.

29.5 Numerical Results

The theory can successfully be applied to a wide variety of problems, such as the hedging of a strip of seasonally options in the gas market where there are 6 underlyings as mentioned in [Section 29.1](#) or the hedging of an option on peak power for which one would execute the hedging with the base product. In the first example, one has to consider 6 different options, each with its own expiry, its own underlying price and hence each with its own hedging ratio. The underlying quantity one would hedge with, would be the average of the 6 underlying contracts (29.1), i.e. the season contract.

The critical timepoints in this framework are the points in time where the underlying monthly forwards become tradeable, the moment where the liquidly traded average forward expires as well as the expiries of the individual options contracts in the strip. In practice, right after the first option expires, the average contract with price F , as well as the first monthly forward with price $F^{(1)}$ will expire. The market then cascades into having a liquid universe of 2 monthly forwards $F^{(2)}$ and $F^{(3)}$, as well as a quarterly forward that consists of the (weighted) average of the last 3 monthly contracts in the strip.

Although perfectly sensible, the numerical example will require the determination of more variables, in particular the 6 volatilities and 15 pair-wise correlations. Moreover, the correlation structure for such gas contracts typically changes as each contract is about to expire because the weather forecasts have a 2-week horizon and the impact on price for near-to delivery forwards is crucial. One can then observe a significant decorrelation of this contract. The LRM framework still applies and is in fact a very powerful technique.

However, the numerical illustration is a lot simpler and clearer when we focus on the electricity forwards where base and peak contracts are for delivery in the same period and do not have different expiries. We restrict ourselves to a setting in which the claim is depending on $F^{(1)}$ but the hedging will be done with $F = w^{(1)}F^{(1)} + w^{(2)}F^{(2)}$. We investigate the hedging for an at-the-money call on $F^{(1)}$ as this is the most challenging example. It is well known that hedging of far in-the-money or out-of-the-money options is relatively easy because the rebalancing frequency is low. The dependency of such options on the underlying forward is either complete (in-the-money) or weak (out-of-the-money). The change per time-step of the required hedging ratios is hence very small.

Note that both $F^{(1)}$ and $F^{(2)}$ are contracts for delivery over different periods of time, while the delivery period of F covers both periods. The total premiums should always be adjusted to the delivery period. We take the example of the base/peak problem, see [Section 29.1](#). This means that the weights are roughly speaking $w^{(1)} = 1/3$ for peak and $w^{(2)} = 2/3$ for off-peak power. The cash flow corresponding to a purchase or sale of such a contract is $w^{(i)}F^{(i)}$ to adjust correctly for the delivery period.

Hence, it is clear that buying the base contract F delivers power during the peak and off-peak hours, corresponding to a position in both assets $F^{(1)}$ and $F^{(2)}$. This means that there are two intuitively choices for hedging the claim on peak power $F^{(1)}$ by using base. One could try and focus on the volume risk or on the price risk, see [Section 29.4.1](#).

29.5.1 Set-Up

We introduce the following notations: $\Gamma(F^{(1)})$ stands for the price of the claim while $\Delta^{(1)} = \frac{\partial \Gamma}{\partial F^{(1)}}$ and $\Delta^{(2)} = \frac{\partial \Gamma}{\partial F^{(2)}}$ represent its partial derivatives with respect to the peak and off-peak contract prices.

Note that in our example the option only depends on $F^{(1)}$ and thus $\Delta^{(2)} = 0$. The amount of risky assets, that are used as a hedge for the claim, is denoted by ξ . For convenience, we will assume that the interest rate $r = 0$. We will observe the expected value and the standard deviation of the total cost, which is defined as the sum of the initial cost and the final deviation of the hedging portfolio from the option in case of a self-financing portfolio. Hence, the L^2 -hedging error, which is often calculated in literature, equals the sum of the variance of the total cost with the square of the difference between the mean of total cost and the initial cost. For a non-self-financing portfolio the total cost is defined as the sum of all the costs on every rebalance date. Given the specific nature of this problem, we assume a lifetime of the claim of $T = 3$ years, where for the first $T_1 = 0.5$ year, a strategy in the base asset is followed. This is inspired by the fact that at some point, liquidity grows in the peak contract. We call the time T_1 the roll-over point. We assume that after this time, the claim can be hedged further with a classical delta hedge or any other hedging strategy on the asset $F^{(1)}$ itself. In a Brownian setting from this point onwards we will hedge perfectly and there is no need in an analysis beyond this point.

The price of the claim at time zero is such that the expected total cost of the strategy is zero, where the price of the option at roll-over time is determined under the unique martingale measure in the Brownian motion case, while in the discontinuous case the mean-correcting martingale measure is used. Due to the zero interest rate, we can restrict ourselves to observing the total cost over the lifetime at the roll-over time. This cost of hedging will be neutralized by the initial premium of the claim. We will show that the uncertainty over the outcome of the different strategies is quite large and therefore we will also study the standard deviation of the hedging cost in those different strategies. The one with the lowest variance is clearly to be preferred in practice.

Both the peak and the off-peak contracts are assumed to follow a geometric Brownian motion:

$$F_t^{(i)} = F_0^{(i)} \exp(\mu^{(i)}t + \sigma^{(i)}W_t^{(i)}), \quad i = 1, 2$$

where the correlation between the Brownian motions is given by $d\langle W^{(1)}, W^{(2)} \rangle_t = \rho dt$.

As parameters we choose $\sigma^{(1)} = 40\%$, $\sigma^{(2)} = 30\%$ and $\rho = 75\%$. For the drift we look at two different situations. The first and most easy one is where we assume both assets to be martingales. Hence $\mu^{(i)} = r - 0.5(\sigma^{(i)})^2$ and we will call this the martingale case. In a second example, the semimartingale case, we introduce a drift by setting $\mu^{(1)} = 0.07$ and $\mu^{(2)} = 0.05$. For both cases, we will look at the performance of the different strategies.

As starting levels for the prices, we assume that $F^{(1)} = \text{€}90/\text{MWh}$ and $F^{(2)} = \text{€}60/\text{MWh}$, and hence the base asset is worth $F = \text{€}70/\text{MWh}$. If we normalize the time of the base contract to one, the cash flows would be given by $\text{€}70$ for baseload of which $\text{€}30$ is coming from the peak contract and $\text{€}40$ from the off-peak. Note that although the price for off-peak is lower, the total cash flow of the off-peak power is higher compared to the one of the peak contract because the amount of delivered hours during off-peak is higher.

29.5.2 Different Strategies

In this section, we repeat the strategies described in [Section 29.4.1](#) and [Section 29.4.2](#), but adjusted to the setting described here, namely where the claim only depends on $F^{(1)}$. The control strategy, which we cannot follow in practice, is added.

29.5.2.1 Control Strategy

In order to verify our results, we calculate the classical strategy. This means that we are hedging the claim on $F^{(1)}$ by effectively taking positions in this asset.

29.5.2.2 Total Volume-Neutral Strategy

The number ξ is here equal to

$$\xi = \max(\Delta^{(1)}, \Delta^{(2)}) = \Delta^{(1)}.$$

We basically focus on the total volume of the peak contract. If the derivative of the claim Γ with respect to $F^{(1)}$ requires a certain amount in $F^{(1)}$, this same amount is taken in F , ensuring that the volume during the peak hours is correct. However, the residual risk that comes into the picture, is the volume taken in the off-peak asset.

In fact, the volatility of the off-peak asset is lower, and therefore ignoring this asset is safe. Clearly, the risk in this strategy is coming mostly from the second risky asset $F^{(2)}$.

29.5.2.3 Volume-Neutral Strategy

The optimal amount ξ is given by

$$\xi = w^{(1)}\Delta^{(1)} + w^{(2)}\Delta^{(2)} = w^{(1)}\Delta^{(1)}.$$

In this strategy, it is assumed that if we need 3 MW of peak power, one can replace this by 1 MW of base power, because the total amount of power over the delivery period is then roughly the same. Or in other words, it is assumed that we can replace volume in the peak hours by volume in the off-peak hours.

It will become clear that this is the worst strategy.

29.5.2.4 Price-Adjusted Strategy

If we take

$$\xi = \frac{w^{(1)}\Delta^{(1)}F^{(1)} + w^{(2)}\Delta^{(2)}F^{(2)}}{w^{(1)}F^{(1)} + w^{(2)}F^{(2)}} = \frac{w^{(1)}\Delta^{(1)}F^{(1)}}{F},$$

the value of the hedge in F and the value of the (theoretical) hedge in $F^{(1)}$ are equal. This ensures that the cash flows during the hedging strategy are the ones one would have from the hedging strategy in $F^{(1)}$.

29.5.2.5 Risky Strategy

If one wants to fully understand the concept of hedging, one should always be prepared to take one step back and ask oneself if the riskiness really decreased by setting up a strategy. Therefore, we compare the strategies to the strategy of doing nothing and waiting until the roll-over time before starting to hedge the claim. In this case, the full risk is taken and $\xi = 0$.

29.5.2.6 Adjusted LRM Hedging Strategy

In none of the above strategies, the volatility or correlation between $F^{(1)}$ and $F^{(2)}$ played a role. It is however very natural that this should have an effect on the strategy one should follow. The (adjusted) LRM

strategy captures this completely, see [Section 29.4.2](#) and in fact outperforms all of the above strategies. In this setting where $\Delta^{(2)} = 0$, the optimal amount equals

$$\xi = \frac{\sigma^{(1)} F^{(1)} \Delta^{(1)} (w^{(1)} \sigma^{(1)} F^{(1)} + w^{(2)} \sigma^{(2)} F^{(2)} \rho)}{w^{(1)} (\sigma^{(1)} F^{(1)})^2 + 2w^{(1)} w^{(2)} \rho \sigma^{(1)} F^{(1)} \sigma^{(2)} F^{(2)} + w^{(2)} (\sigma^{(2)} F^{(2)})^2}.$$

29.5.3 Results

To obtain the results, we simulated 25 000 paths and rebalance twice a week.

29.5.3.1 Hedging Cost

For the various strategies, we determine the hedging cost up to the roll-over time. We look at this cost both for the martingale case as well as for the semimartingale case. Table 29.1 contains for each case the expected cost and the standard deviation between brackets. The larger this standard deviation, the more uncertainty and hence the more risk remains in the hedging procedure. Let us focus first on the martingale case. It is obvious that the LRM strategy outperforms the current market practices. Compared to doing nothing, hence this is the full risk case, the improvement is very good. The reason that the total volume-neutral strategy works well is because a big part of the risk is concentrated in the peak price since this contract has the highest volatility.

In practice, the hedging cost is considered as the fair value price of the option. From Table 29.1, we deduce that the average cost of hedging is almost identical across all the strategies, hence each strategy indicates the same fair value price for the option at the start.

In the semimartingale case, we can observe that for the control strategy, there is no effect. This is natural as we already know that pricing is always done under a risk-neutral martingale measure, which is unique in the continuous case. For all the other strategies, we see that the cost of the strategy is changing. At the same time, the uncertainty grows as well. However, once again the LRM behaves better than any of the others.

Remark that in practice, the estimation of the drift term is virtually impossible. Knowing the drift would mean, knowing where the prices would go and often it might be possible to distinguish trends in the short or extremely long run, but the deviations from these kind of trends make it very hard to even estimate the drift term correctly. Since the energy market is a forward market, we can assume that the market prices everything correctly, and hence that the quantities are indeed martingales. If later, it turns out that there was a systematic drift, we then hope that the margin taken at inception in the option premium is sufficient to cover this.

We conclude that the LRM hedging strategy outperforms the more intuitive approaches and even in case the assets are only semimartingales, the method still works well.

In fact, we could even go one step further. We want to calculate the cost of hedging in case the underlyings follow a discontinuous price process, but in a very fast way and hence by avoiding again the use

TABLE 29.1 Hedging Cost and Standard Deviation in the Case of Brownian Motions

Strategy	Martingale case	Semimartingale case
Control	24.39 (0.39)	24.39 (0.40)
Total volume-neutral	24.37 (8.24)	26.06 (9.30)
Volume-neutral	24.34 (13.94)	27.84 (15.77)
Price-adjusted	24.35 (12.95)	27.57 (14.59)
(Adjusted) LRM	24.38 (5.95)	24.67 (6.33)
Full risk	24.33 (17.15)	28.73 (19.38)

of PIDE's. Therefore, we use the amount ξ described in (29.10), where the price of the claim and the Δ 's are calculated under the mean-correcting martingale measure linked with the two processes $F^{(1)}$ and $F^{(2)}$. For this purpose, we assume that both peak and off-peak can be written as exponential variance gamma processes, where the characteristic function of a variance gamma function equals

$$\phi(z) = \left(1 - i u \theta v + \frac{1}{2} \sigma^2 v u^2 \right)^{-1/v},$$

see Schoutens (2003). As in Leoni and Schoutens (2008), we assume that the Gamma clock is equal for both assets and hence they jump at the same time. We will take the following parameters:

$\mu_{\text{mart}}^{(1)} = -0.0179$	$\mu_{\text{mart}}^{(2)} = -0.005$
$\mu_{\text{semi}}^{(1)} = 0.07$	$\mu_{\text{semi}}^{(2)} = 0.05$
$\sigma^{(1)} = 40.50\%$	$\sigma^{(2)} = 30\%$
$\theta^{(1)} = -0.10$	$\theta^{(2)} = -0.05$
$v = 0.25$	
$\rho = 74.80\%$	

These numbers ensure us that the option price, calculated under the mean-correcting measure for peak, leads to the same price as we had in the Brownian case. The correlation between the Brownian components has been adjusted downwards such that the linear correlation coefficient between the log-returns of the assets remained around 75% as earlier. Furthermore, $\mu^{(i)}$ in the martingale case is determined by the following equation:

$$\mu_{\text{mart}}^{(i)} = r - \log(1 - 0.5v(\sigma^{(i)})^2 - \theta^{(i)}v) / v, \quad i = 1, 2.$$

Within this set-up, we obtain the results reported in Table 29.2. In the martingale case, all the strategies have a lower cost of hedging, but with a greater uncertainty than in the Brownian motion case. This can be explained by the fact that within a VG model, there are only small changes in the prices of the assets until a significant jump is noticeable. The fat-tailed distribution (compared to the normal distribution) favours smaller moves most of the time and some extreme jumps once in a while.

The interesting aspect of this analysis is that the control strategy becomes less good in the sense that the uncertainty becomes bigger in the semimartingale case. This can be explained easily by noting that the introduction of jumps induce a higher variance in the hedging process if the rebalancing frequency remains equal.

We can also observe that the impact, at least relatively speaking on the other strategies of the introduction of such jumps is less important and the uncertainty remains completely driven by the basis risk while hedging. Turning this around, one can also understand that increasing the hedging frequency will not eliminate the variance in the RM strategy as one cannot expect the error to become infinitely small

TABLE 29.2 Hedging Cost and Standard Deviation in the Case of a Multivariate Variance Gamma Process

Strategy	Martingale case	Semimartingale case
Control	24.39 (2.41)	24.39 (2.43)
Total volume-neutral	23.88 (9.08)	24.81 (9.49)
Volume-neutral	23.63 (14.74)	25.54 (15.40)
Price-adjusted	23.67 (13.74)	25.44 (14.33)
(Adjusted) LRM	24.06 (6.65)	24.25 (6.97)
Full risk	23.51 (17.94)	25.91 (18.75)

by just hedging more often. The variance or risk is introduced by the liquidity constraint and the fact that one chooses to hedge with a correlated, but different asset.

When we turn to the semimartingale case, we can deduce similar results as before. The cost of hedging depends on the actual drift of the process and in general this is not a nice feature of a strategy because this drift is extremely hard to measure or estimate. However, it becomes clear in this case as well, that the LRM strategy is rather robust, making it the most suitable candidate for real hedging of claims on non-tradable assets.

In fact, one could also argue that we do not know what the real LRM hedging strategy would do and if this would not behave even better. We think the real strategy will surely not perform better than in the continuous martingale case, and our adjusted LRM hedging strategy only performs slightly worse than the exact LRM hedging strategy used in the continuous martingale case. Hence, we believe that the possible increase in accuracy of the results will not outweigh for the loss in computational speed. Furthermore, the reason why this adjusted strategy works well, even in the discontinuous case, is because it still accurately captures the real correlation between the two assets, although it does not use the exact angle bracket process. The correct LRM however, would require a transformed drift process, which would shift the expected cost of hedging (and hence the price of the claim), but the effect on the variance would be smaller. For small drift terms, the approximation works well enough for practical purposes.

It should be clear that the LRM theory can be applied to any complex process and not just a variance gamma driven process. Typically, more complex processes can provide better models but the number of parameters will increase. Moreover, the variance gamma process has the convenient property that the covariance terms can be explicitly written down, see Leoni and Schoutens (2008), eliminating a numerical calibration.

29.6 Conclusion

The focus in this paper was to find a numerically applicable and fast method to hedge options on non-traded assets. Thereto, we first derived theoretically the formulas for the LRM hedging strategy for this non-traded asset case. This theoretical result motivated the definition of the adjusted LRM hedging strategy as an approximation to the exact one. In the numerical part, we showed that this adjusted strategy outperforms the current market practices in the Brownian motion case as well as in the variance gamma case.

Acknowledgements

Nele Vandaele gratefully acknowledges the financial support by the Research Foundation-Flanders. The authors thank two anonymous referees and the editor for the valuable suggestions that improved the paper considerably.

References

- Albert, A., *Regression and the Moore–Penrose Pseudoinverse, Volume 94 of Mathematics in Science and Engineering*, 1972 (Academic Press: New York and London).
- Altmann, T., Schmidt, T. and Stute, W., A shot noise model for financial assets. *Int. J. Theor. Appl. Finance*, 2008, **11**(1), 87–106.
- Ankirchner, S., Imkeller, P. and Reis, G., Pricing and hedging of derivatives based on non-tradable underlyings. *Math. Finance*, 2010, **20**(2), 289–312.
- Bertsimas, D., Kogan, L. and Lo, A.W., Hedging derivative securities and incomplete markets: An ϵ -arbitrage approach. *Oper. Res.*, 2001, **49**(3), 372–397.

- Biagini, F., Guasoni, P. and Pratelli, M., Mean-variance hedging for stochastic volatility models. *Math. Finance*, 2000, **10**(2), 109–123.
- Brodén, M. and Tankov, P., Tracking errors from discrete hedging in exponential Lévy models. *Int. J. Theor. Appl. Finance*, 2011, **14**, 1–35.
- Černý, A. and Kallsen, J., On the structure of general mean-variance hedging strategies. *Ann. Probab.*, 2007, **35**(4), 1479–1531.
- Černý, A. and Kallsen, J., Mean-variance hedging and optimal investment in Heston's model with correlation. *Math. Finance*, 2008, **18**, 473–492.
- Chan, T., Kollar, J. and Wiese, A., The variance-optimal martingale measure in Lévy models with stochastic volatility. Technical report, 2009.
- Choulli, T., Krawczyk, L. and Stricker, C., ε -martingales and their applications in mathematical finance. *Ann. Probab.*, 1998, **26**(2), 853–876.
- Choulli, T., Vandaele, N. and Vanmaele, M., The Föllmer–Schweizer decomposition: Comparison and description. *Stochas. Process. Appl.*, 2010, **120**(6), 853–872.
- Colwell, D. and Elliott, R., Discontinuous asset prices and non-attainable contingent claims. *Math. Finance*, 1993, **3**(3), 295–308.
- Davis, M., Optimal hedging with basis risk. In *From Stochastic Calculus to Mathematical Finance, Mathematics and Statistics*, edited by Y. Kabanov, R. Liptser, and J. Stoyanov, pp. 169–187, 2006 (Springer: Berlin Heidelberg).
- Denkl, S., Goy, M., Kallsen, J., Muhle-Karbe, J. and Pauwels, A., On the performance of delta hedging strategies in exponential Lévy models. Technical report, Universität zu Kiel, 2009.
- El Karoui, N., Peng, S. and Quenez, M., Backward stochastic differential equations in finance. *Math. Finance*, 1997, **7**(1), 1–71.
- Frey, R. and Runggaldier, W., Risk-minimizing hedging strategies under restricted information: The case of stochastic volatility models observable only at discrete random times. *Math. Methods Oper. Res.*, 1999, **50**, 339–350.
- Henderson, V., Valuation of claims on nontraded assets using utility maximization. *Math. Finance*, 2002, **12**(4), 351–373.
- Henderson, V. and Hobson, D., Real options with constant relative risk aversion. *J. Econ. Dyn. Control*, 2002, **27**, 329–355.
- Hobson, D.G., Bounds for the utility-difference prices of non-traded assets in incomplete markets. *Decis. Econ. Finance*, 2005, **28**, 33–52.
- Horst, U., Pirvu, T. and Reis, G., On securitization, market completion and equilibrium risk transfer. *Math. Financ. Econ.*, 2010, **2**(4), 211–252.
- Hubalek, F. and Sgarra, C., Quadratic hedging for the Bates model. *Int. J. Theor. Appl. Finance*, 2007, **10**(5), 873–885.
- Hubalek, F., Kallsen, J. and Krawczyk, L., Variance-optimal hedging for processes with stationary independent increments. *Ann. Appl. Probab.*, 2006, **16**, 853–885.
- Kallsen, J. and Pauwels, A., Variance-optimal hedging in general affine stochastic volatility. *Adv. Appl. Probab.*, 2010, **42**(1), 83–105.
- Kallsen, J. and Pauwels, A., Variance-optimal hedging for time changed Lévy processes. *Appl. Math. Finance*, 2011, **18**(1), 1–28.
- Kallsen, J. and Vierthauer, R., Quadratic hedging in affine stochastic volatility models. *Rev. Deriv. Res.*, 2009, **12**(1), 3–27.
- Kallsen, J., Muhle-Karbe, J., Shenkman, N. and Vierthauer, R., Discrete-time variance-optimal hedging in affine stochastic volatility models. In *Alternative Investments and Strategies*, edited by R. Kiesel, M. Scherer, and R. Zagst, pp. 369–387, 2010 (World Scientific: Singapore).
- Leoni, P. and Schoutens, W., Multivariate smiling, *Wilmott Mag.*, March 2008.
- Monat, P. and Stricker, C., Föllmer–Schweizer decomposition and mean-variance hedging for general claims. *Ann. Probab.*, 1995, **23**, 605–628.

- Monoyios, M., Performance of utility-based strategies for hedging basis risk. *Quant. Finance*, 2004, **4**, 245–255.
- Njoh, S., Cross hedging within a log mean reverting model. *Int. J. Theor. Appl. Finance*, 2007, **10**(5), 887–914.
- Pham, H., On quadratic hedging in continuous time. *Math. Methods Oper. Res.*, 2000, **51**, 315–339.
- Poulsen, R., Schenk-Hoppé, K.R. and Ewald, C.-O., Risk minimization in stochastic volatility models: Model risk and empirical performance. *Quant. Finance*, 2009, **9**(6), 693–704.
- Schoutens, W., *Lévy Processes in Finance*. Wiley Series in Probability and Statistics, 2003 (John Wiley & Sons: Chichester).
- Schweizer, M., Option hedging for semimartingales. *Stochas. Process. Appl.*, 1991, **37**, 339–363.
- Schweizer, M., A guided tour through quadratic hedging approaches. In *Option Pricing, Interest Rates and Risk Management*, edited by E. Jouini, J. Cvitanić, and M. Musiela, pp. 538–574, 2001 (Cambridge University Press: Cambridge, UK).
- Schweizer, M., Local risk-minimization for multidimensional assets and payment streams. *Banach Cent. Pub.*, 2008, **83**, 213–229.
- Vandaele, N., Quadratic hedging in finance and insurance. PhD Thesis, Ghent University, 2010.
- Vandaele, N. and Vanmaele, M., Overview: (locally) risk-minimizing hedging strategies. In *Handelingen Contactforum Actuarial and Financial Mathematics Conference, February 7–8, 2008*, edited by M. Vanmaele, G. Deelstra, A. De Schepper, J. Dhaene, H. Reynaerts, W. Schoutens, and P. Van Goethem, pp. 107–119, 2008 (Koninklijke Vlaamse Academie van België voor Wetenschappen en Kunsten: Brussels).

The Valuation of Clean Spread Options: Linking Electricity, Emissions and Fuels

30.1	Introduction	645
30.2	The Bid Stack: Price Setting in Electricity Markets.....	647
30.3	Risk-Neutral Pricing of Allowance Certificates.....	648
	The Market Emission Rate • The Pricing Problem • An FBSDE for the Allowance Price • Existence of a Solution to the Pricing Problem	
30.4	Valuing Clean Spread Options	652
30.5	A Concrete Two-Fuel Model.....	653
	The Bid Stack • The Emission Stack • Specifying the Exogenous Stochastic Factors	
30.6	Numerical analysis	655
	Choice of parameters • Case Study I: Impact of the Emission Market • Case study II: Impact of fuel price changes • Case Study III: Comparison with Reduced-Form • Case Study IV: Cap-and-Trade vs. Carbon Tax	
30.7	Conclusion.....	662
	Acknowledgements	662
	References.....	662
	Appendix A: Numerical Solution of the FBSDE	663
	Appendix B: Numerical Calculation of Spread Prices.....	665

René Carmona

Michael Coulon

Daniel Schwarz

The purpose of the paper is to present a new pricing method for clean spread options, and to illustrate its main features on a set of numerical examples produced by a dedicated computer code. The novelty of the approach is embedded in the use of a structural model as opposed to reduced-form models which fail to capture properly the fundamental dependencies between the economic factors entering the production process.

Keywords: Spread options; Real asset valuation; Electricity markets; Emission markets

JEL Classification: G1, G10, G12

30.1 Introduction

Spread options are most often used in the commodity and energy markets to encapsulate the profitability of a production process by comparing the price of a refined product to the costs of production, including, but not limited to, the prices of the inputs to the production process. When the output

commodity is electric power, such spread options are called *spark spreads* when the electricity is produced from natural gas, and *dark spreads* when the electricity is produced from coal. Both processes are the sources of Greenhouse Gas (GHG) emissions, in higher quantities for the latter than the former. In this paper we concentrate on the production of electricity and CO₂ emissions and the resulting dependence structure between prices.

Market mechanisms aimed at controlling CO₂ emissions have been implemented throughout the world, and whether they are mandatory or voluntary, cap-and-trade schemes have helped to put a price on carbon in Europe, the US, and around the world. In the academic literature, equilibrium models have been used to show what practitioners have known all along, namely that the price put on CO₂ by the regulation should be included in the costs of production to set the price of electricity (Carmona *et al.* 2010).

Strings of spark spread options (options on the spread between the price of 1 MWh of electricity and the cost of the amount of natural gas needed to produce such a MWh) with maturities covering a given period are most frequently used to value the optionality of a gas power plant which can be run when it is profitable to do so (namely when the price of electricity is greater than the cost of producing it), and shut down otherwise. In a nutshell, if an economic agent takes control on day t of a gas power plant for a period $[T_1, T_2]$, then for every day $\tau \in [T_1, T_2]$ of this period, he or she can decide to run the power plant when $P_\tau > h_g S_\tau^g + K$, booking a profit $P_\tau - h_g S_\tau^g - K$ for each unit of power produced, and shut the plant down if $P_\tau \leq h_g S_\tau^g + K$. Here P_τ denotes the price at which one unit (1 MWh) of power can be sold on day τ , S_τ^g the price of one unit of natural gas (typically one MMBtu), h_g the efficiency or heat rate of the plant (i.e. the number of units of natural gas needed to produce one unit of electricity) and K the daily (fixed) costs of operations and maintenance of the plant. Ignoring constraints such as ramp-up rates and start-up costs, this scheduling is also automatically induced when generators bid at the level of their production costs in the day-ahead auction for power. So in this somewhat oversimplified analysis of the optionality of the plant, the value at time t of the control of the plant operation on day τ can be expressed as $e^{-r(\tau-t)} \mathbb{E}[(P_\tau - h_g S_\tau^g - K)^+ | \mathcal{F}_t]$, where, as usual, the exponent $+$ stands for the positive part, i.e. $x^+ = x$ when $x \geq 0$ and $x^+ = 0$ otherwise, r for the constant interest rate used as discount factor to compute the present values of future cash flows, and $\mathcal{F}_t$ denotes the information available on day t when the conditional expectation is actually computed. So the operational control (for example, as afforded by a tolling contract) of the plant over the period $[T_1, T_2]$ could be valued on day t as

$$V_t^{PP} = \sum_{\tau=T_1}^{T_2} e^{-r(\tau-t)} \mathbb{E}[(P_\tau - h_g S_\tau^g - K)^+ | \mathcal{F}_t].$$

This rather simplistic way of valuing a power generation asset in the spirit of the theory of real options had far-reaching implications in the developments of the energy markets, and is the main reason why spread options are of the utmost importance. However, such a valuation procedure is flawed in the presence of emission regulation as the costs of production also have to include the costs specific to the regulation. To be more specific, the day- τ potential profit $(P_\tau - h_g S_\tau^g - K)^+$ of the spark spread has to be modified to $(P_\tau - h_g S_\tau^g - e_g A_\tau - K)^+$ in order to accommodate the cost of the regulation. Here A_τ is the price of one allowance certificate worth one ton of CO₂ equivalent, and e_g is the emission coefficient of the plant, namely the number of tons of CO₂ emitted by the plant during the production of one unit of electricity. Such a spread is often called a *clean spread* to emphasize the fact that the externality is being paid for, and the real option approach to power plant valuation leads to the following *clean price*:

$$V_t^{CPP} = \sum_{\tau=T_1}^{T_2} e^{-r(\tau-t)} \mathbb{E}[(P_\tau - h_g S_\tau^g - e_g A_\tau - K)^+ | \mathcal{F}_t]$$

for the control of the plant over the period $[T_1, T_2]$ in the presence of the regulation.

In order to price such cross-commodity derivatives, a joint model is clearly required for fuel prices, electricity prices and carbon allowance prices. Various studies have analysed the strong links between these price series (De Jong and Schneider 2009, Koenig 2011). Many reduced-form price models have

been proposed for electricity (see Eydeland and Wolyniec (2003) and Benth *et al.* (2008) reviews) with a focus on capturing its stylized features such as seasonality, high volatility, spikes, mean-reversion and fuel price correlation. On the other hand, many authors have argued that these same features are better captured via a structural approach, modelling the dynamics of underlying factors such as demand (load), capacity and fuel prices (early examples include Barlow (2002), Cartea and Villaplana (2008), Pirrong and Jermakyan (2008) and Coulon and Howison (2009)).

Similarly, for carbon emission allowances, exogenously specified processes that model prices directly have been proposed by some (Carmona and Hinz 2011). Others have instead treated the emission process as the exogenously specified underlying factor; in this case the allowance certificate becomes a derivative on cumulative emissions (Seifert *et al.* 2008, Chesney and Taschini 2012). However, these models do not take into account the important feedback from the allowance price to the rate at which emissions are produced in the electricity sector—a feature that is crucial for the justification of any implementation of a cap-and-trade scheme. In a discrete-time framework this feedback mechanism has been addressed, for example, by Coulon (2009) and Carmona *et al.* (2010). In continuous time the problem has been treated by Carmona *et al.* (2012b) and Howison and Schwarz (2012), whereby the former models the accumulation of emissions as a function of an exogenously specified electricity price process, while the latter uses the bid stack mechanism to infer the emission rate.

The literature on spread options is extensive. In industry, Margrabe's classical spread option formula (Margrabe 1978) is still widely used, and has been extended by various authors (see Carmona and Durrleman (2003) for an overview), including to the three commodity case, as required for the pricing of clean spreads (Alos *et al.* 2011). Carmona and Sun (2012) analyse the pricing of two-asset spread options in a multiscale stochastic volatility model. For electricity markets, pricing formulae for dirty spreads based on structural models have been proposed by Carmona *et al.* (2012a), who derive a closed-form formula in the case of $K = 0$, and by Aïd *et al.* (2012), who derive semi-closed form formulae for $K \neq 0$ at the expense of a fixed merit order.

The original contributions of the paper are twofold. First, we express the value of clean spread options in a formulation where demand for power and fuel prices are the only factors whose stochastic dynamics are given exogenously, and where the prices of power and emission allowances are derived from a bid-stack-based structural model and a forward backward stochastic differential system, respectively. The second contribution is the development of a numerical code for the computation of the solution of the pricing problem. First, we solve a 4 + 1-dimensional semilinear partial differential equation to compute the price of an emission allowance, then we use Monte Carlo techniques to compute the price of the spread option. These computational tools are used to produce the numerical results for the case studies presented in Section 30.6 of the paper for the purpose of illustrating the impact of a carbon regulation on the price of spread options. In this final section we first compare the price of spark and dark spread options in two different markets, one with no emission regulation and the other governed by an increasingly strict cap-and-trade system. Second, we analyse the impact that different merit order scenarios have on the option prices. Third, we demonstrate the difference between the structural and the reduced-form approach by comparing the option prices produced by our model with those produced by two candidate reduced-form models. Fourth and last, we contrast two competing policy instruments: cap-and-trade, represented by the model we propose, and a fixed carbon tax.

30.2 The Bid Stack: Price Setting in Electricity Markets

In order to capture the dependency of the electricity price on production costs and fundamental factors in a realistic manner, we use a structural model in the spirit of those reviewed in the recent survey of Carmona and Coulon (2012). The premises of structural models for electricity prices depend upon an explicit construction of the supply curve. Since electricity is sold at its marginal cost, the electricity spot price is given by evaluating the supply function for the appropriate values of factors used to describe the costs of production in the model.

In practice, electricity producers submit day-ahead bids to a central market operator, whose task it is to allocate the production of electricity amongst them. Typically, firms' bids have the form of price–quantity pairs, with each pair comprising the amount of electricity the firm is willing to produce, and the price at which the firm is willing to sell this quantity. Given the large number of generators in most markets, it is common in structural models to approximate the resulting step function of market bids by a continuous increasing curve. Firms' bid levels are determined by their costs of production. An important feature of our model, distinguishing it from most of the commonly used structural models, is the inclusion, as part of the production costs, of the costs incurred because of the existence of emission regulation.

We assume that, when deciding which firms to call upon to produce electricity, the market operator adheres to the merit order, a rule by which cheaper production units are called upon before more expensive ones. For simplicity, operational and transmission constraints are not considered.

Assumption 30.2.1

The market operator arranges bids according to the merit order, in increasing order of production costs.

The map resulting from ordering market supply in increasing order of electricity costs of production is what is called the bid stack. As it is one of the important building blocks of our model, we define it in a formal way for later convenience.

Definition 30.2.2: The *bid stack* is given by a measurable function

$$b : [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^n \ni (x, a, s) \mapsto b(x, a, s) \in \mathbb{R},$$

with the property that, for each fixed $(a, s) \in \mathbb{R} \times \mathbb{R}^n$, the function $[0, \bar{x}] \ni x \mapsto b(x, a, s)$ is strictly increasing.

In this definition, $\bar{x} \in \mathbb{R}_{++}$ represents the *market capacity* (measured in MWh) and the variable x the *supply of electricity*. The integer $n \in \mathbb{N} \setminus \{0\}$ gives the number of economic factors (typically the prices, in € say, of the fuels used in the production of electricity), and $s \in \mathbb{R}^n$ the numeric values of these factors. Here and throughout the rest of the paper the *cost of carbon emissions* (measured in € per metric ton of CO₂) is denoted by a . So for a given allowance price, say a , and fuel prices, say s , the market is able to supply x units of electricity at price level $b = b(x, a, s)$ (measured in € per MWh). In other words, $b(x, a, s)$ represents the bid level of the marginal production unit in the event that demand equals x .

The choice of a function b which captures the subtle dependence of the electricity price upon the level of supply and the production costs, is far from trivial, and different approaches have been considered in the literature, as reviewed recently by Carmona and Coulon (2012). In [Section 30.5.1](#) we extend the model proposed by Carmona *et al.* (2012a) to include the cost of carbon as part of the variable costs driving bid levels.

30.3 Risk-Neutral Pricing of Allowance Certificates

As the inclusion of the cost of emission regulation in the valuation of spread options is the main thrust of the paper, we explain how emission allowances are priced in our model. The model we introduce is close to Howison and Schwarz (2012). However, we extend the results found therein to allow the equilibrium bids of generators to be stochastic and driven by fuel prices, a generalization that is vital for our purpose.

We assume that carbon emissions in the economy are subject to cap-and-trade regulation structured as follows: at the end of the compliance period, each registered firm needs to offset its cumulative emissions with emission allowances or incur a penalty for each excess ton of CO₂ not covered by a redeemed allowance certificate. Initially, firms acquire allowance certificates through free allocation, e.g. through National Allocation Plans (NAP) like in the initial phase of the European Union (EU) Emissions Trading Scheme (ETS), or by purchasing them at auctions as in the Regional Greenhouse Gas Initiative (RGGI) in the North East of the US. Allowances change hands throughout the compliance period. Typically, a firm which thinks that its initial endowment will not suffice to cover its

emissions will buy allowances, while firms expecting a surplus will sell them. Adding to these *naturals*, speculators enter the market providing liquidity. Allowances are typically traded in the form of forward contracts and options. In this paper, we denote by A_t the spot price of an allowance certificate maturing at the end of the compliance period. Because their cost of carry is negligible, we treat them as financial products liquidly traded in a market without frictions, and in which long and short positions can be taken.

In a competitive equilibrium, the level of cumulative emissions relative to the cap (i.e. the number of allowance certificates issued by the regulation authority) determines whether—at the end of the compliance period—firms will be subjected to a penalty payment and create a demand for allowance certificates (see Carmona *et al.* (2010) for details). For this reason, allowance certificates should be regarded as derivatives on the emissions accumulated throughout the regulation period. This type of option written on a non-tradable underlying interest is rather frequent in the energy markets: temperature options are a case in point.

30.3.1 The Market Emission Rate

As evidenced by the above discussion, the rate at which CO₂ is emitted into the atmosphere as a result of electricity production has to be another important building block of our model. Clearly, at any given time, this rate is a function of the amount of electricity produced and because of their impact on the merit order, the variable costs of production, including fuel prices, and, notably, the carbon allowance price itself.

Definition 30.3.1: The *market emission rate* is given by a measurable function

$$\mu_e : [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^n \ni (x, a, s) \mapsto \mu_e(x, a, s) \in \mathbb{R}_+,$$

on which we will impose technical assumptions later on.

With the above definition, for a given level of electricity supply and for given allowance and fuel prices, $\mu_e = \mu_e(x, a, s)$ represents the rate at which the market emits, measured in tons of CO₂ per hour. Cumulative emissions are then computed by integrating the market emission rate over time. Since any increase in supply can only increase the emission rate, it is of course reasonable from a modelling point of view to expect μ_e to be increasing as a function of x . Similarly, as the cost of carbon increases, the variable costs (and hence the bids) of pollution-intensive generators increase by more than those of environmentally friendlier ones. Dirtier technologies become relatively more expensive and are likely to be scheduled further down in the merit order. As a result, cleaner technologies are brought online earlier and we expect μ_e to be decreasing as a function of a .

In [Section 30.5.2](#) we propose a specific functional form for μ_e consistent with the bid stack model introduced in [Section 30.5.1](#).

30.3.2 The Pricing Problem

We shall use the following notation. For a fixed time horizon $T \in \mathbb{R}_+$, let $(W_t^0, W_t)_{t \in [0, T]}$ be* a $(n + 1)$ -dimensional standard Wiener process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $\mathcal{F}^0 := (\mathcal{F}_t^0)$ the filtration generated by W^0 , $\mathcal{F}^W := (\mathcal{F}_t^W)$ the filtration generated by W , and $\mathcal{F} := \mathcal{F}^0 \vee \mathcal{F}^W$ the market filtration. All relationships between random variables are to be understood in the almost sure sense.

Consumers' *demand for electricity* is given by an $\mathcal{F}^0$ -adapted stochastic process (D_t) taking values in $[0, \bar{x}]$. In response to this demand, producers supply electricity, and we assume that demand and supply are always in equilibrium, so that at any time $t \in [0, T]$ an amount D_t of electricity is supplied. The *prices of fuels* are observed $\mathcal{F}^W$ -adapted stochastic processes (S_t) taking values in $\mathbb{R}^n$, where $S_t := (S_t^1, \dots, S_t^n)$. As

* To simplify the presentation, from now on we drop the subscript $t \in [0, T]$, which specifies the time interval on which a stochastic process is defined.

we will see in Section 30.3.3, the price of an *allowance certificate* at time t , say A_t , is now constructed as a $\mathcal{F}$ -adapted stochastic process solving a Forward Backward Stochastic Differential Equation (FBSDE). The rate of emission $\mu_e(D_t, A_t, S_t)$ can then be evaluated and the cumulative emissions computed by integrating over time, resulting in a $\mathcal{F}$ -adapted process (E_t) .

Since we do not present a calibration of the model to any particular electricity or emissions market, we avoid the difficulties of estimating market prices of risk (see, for example, Eydeland and Wolyniec (2003) and Carmona and Coulon (2012) for discussions of some possible ways to approach this thorny issue), and instead choose to specify the dynamics of the processes (D_t) and (S_t) directly under a risk-neutral measure $\mathbb{Q} \sim \mathbb{P}$ chosen by the market for pricing purposes. In practice, various alternative approaches to parameter estimation and calibration could be used to identify a risk-neutral measure which is consistent with liquidly traded products such as power and fuel forward contracts, thereby also estimating the magnitude of risk premia. However, our focus here is on the overall nature of energy price correlations, their structural origins and their important impact on spread option prices, rather than carrying out a study of any one market. For this reason, we also ignore any questions of market incompleteness (i.e. the uniqueness of $\mathbb{Q}$), transaction costs, illiquidity, inelasticity of demand (and also any risk that demand could ever exceed supply, for example causing a black-out), and finally the role of non-power-sector emissions. While interesting details of the markets, we argue that the inclusion of such effects should not cause any substantial change to the important qualitative conclusions drawn from the model in Section 30.6.

30.3.3 An FBSDE for the Allowance Price

We assume that, at time $t = 0$, demand for electricity is known. Thereafter, it evolves according to an Itô diffusion. Specifically, for $t \in [0, T]$, demand for electricity D_t is the unique strong solution of a stochastic differential equation of the form

$$dD_t = \mu_d(t, D_t)dt + \sigma_d(D_t)d\tilde{W}_t^0, \quad D_0 = d_0 \in (0, \bar{x}), \quad (30.1)$$

where $(\tilde{W}_t^0)$ is an $\mathcal{F}_t^0$ -adapted $\mathbb{Q}$ -Brownian motion. The time dependence of the drift allows us to capture the seasonality observed in electricity demand, and the resulting seasonality in prices.

Similarly to demand, the prices of the fuels used in the production process satisfy a system of stochastic differential equations written in vector form as follows:

$$dS_t = \mu_s(S_t)dt + \sigma_s(S_t)d\tilde{W}_t, \quad S_0 = s_0 \in \mathbb{R}^n, \quad t \in [0, T], \quad (30.2)$$

where $(\tilde{W}_t)$ is an $\mathcal{F}_t^W$ -adapted $\mathbb{Q}$ -Brownian motion. We note that, in some cases, it may be appropriate to also include time-dependence in the drift or volatility above, in order to capture seasonal patterns in some fuels such as natural gas.

Cumulative emissions are measured from the beginning of the compliance period at time $t = 0$, so that $E_0 = 0$. Subsequently, they are determined by integrating the market emission rate μ_e introduced in definition 30.3.1. So assuming that the price A_t of an allowance certificate is known, the cumulative emissions process is represented by an absolutely continuous process, i.e. for $t \in [0, T]$,

$$dE_t = \mu_e(D_t, A_t, S_t)dt, \quad E_0 = 0. \quad (30.3)$$

Note that, with this definition, the process (E_t) is non-decreasing, which makes intuitive sense considering that it represents a cumulative quantity.

To complete the formulation of the pricing model, it remains to characterize the allowance certificate price process $(A_t)_{t \in [0, T]}$. If our model is to apply to a one compliance period scheme, in a competitive equilibrium, at the end of the compliance period $t = T$, its value is given by a deterministic function of the cumulative emissions:

$$A_T = \phi(E_T), \quad (30.4)$$

where $\phi: \mathbb{R} \hookrightarrow \mathbb{R}$ is bounded, measurable and non-decreasing. Usually, $\phi(\cdot) := \pi_{[\Gamma, \infty)}(\cdot)$, where $\pi \in \mathbb{R}_+$ denotes the penalty paid in the event of non-compliance and $\Gamma \in \mathbb{R}_+$ the cap chosen by the regulator as the aggregate allocation of certificates (see Carmona *et al.* (2010) for details). Since the discounted allowance price is a martingale under $\mathbb{Q}$, it is equal to the conditional expectation of its terminal value, i.e.

$$A_t = \exp(-r(T-t))\mathbb{E}^{\mathbb{Q}}[\phi(E_T) | \mathcal{F}_t], \quad \text{for } t \in [0, T], \quad (30.5)$$

which implies, in particular, that the allowance price process (A_t) is bounded. Since the filtration $(\mathcal{F}_t)$ is being generated by the Wiener process, it is a consequence of the Martingale Representation Theorem (Karatzas and Shreve 1999) that the allowance price can be represented as an Itô integral with respect to the Brownian motion $(\tilde{W}_t^0, \tilde{W}_t)$. It follows that

$$dA_t = rA_t dt + z_t^0 d\tilde{W}_t^0 + Z_t \cdot d\tilde{W}_t, \quad \text{for } t \in [0, T], \quad (30.6)$$

for some $\mathcal{F}_t$ -adapted, square integrable process (Z_t^0, Z_t) .

Combining equations (30.1), (30.2), (30.3), (30.4) and (30.6), the pricing problem can be reformulated as the solution of the FBSDE

$$\begin{cases} dD_t = \mu_d(t, D_t)dt + \sigma_d(D_t)d\tilde{W}_t^0, & D_0 = d_0 \in (0, \bar{x}), \\ dS_t = \mu_s(S_t)dt + \sigma_s(S_t)d\tilde{W}_t, & S_0 = s_0 \in \mathbb{R}^n, \\ dE_t = \mu_e(D_t, A_t, S_t)dt, & E_0 = 0, \\ dA_t = rA_t dt + Z_t^0 d\tilde{W}_t^0 + Z_t \cdot d\tilde{W}_t, & A_T = \phi(E_T). \end{cases} \quad (30.7)$$

Note that the first two equations are standard stochastic differential equations (in the forward direction of time) which do not depend upon the cumulative emissions and the allowance price. We will choose their coefficients so that existence and uniqueness of solutions hold. Also, for the sake of convenience, we implicitly assume that the function μ_e , whose first argument was originally restricted to the interval $[0, \bar{x}]$, is defined on the whole $\mathbb{R} \times \mathbb{R} \times \mathbb{R}^n$ by setting $\mu_e(x, a, s) = \mu_e(0, a, s)$ for $x < 0$ and $\mu_e(x, a, s) = \mu_e(\bar{x}, a, s)$ for $x > \bar{x}$. Finally, we make the following assumptions on the coefficients of (30.7).

Assumption 30.3.2

The functions $\mu_d: [0, T] \times [0, \bar{x}] \hookrightarrow \mathbb{R}$, $\sigma_d: [0, \bar{x}] \hookrightarrow \mathbb{R}$, $\mu_s: \mathbb{R}^n \hookrightarrow \mathbb{R}^n$, $\sigma_s: \mathbb{R}^n \hookrightarrow \mathbb{R}^n \times \mathbb{R}^n$ are such that the first two equations in (30.7) have a unique strong solution.

30.3.4 Existence of a Solution to the Pricing Problem

Theorem 30.3.3:

If Assumption 30.3.2 holds, the function μ_e giving the emission rate is Lipschitz with respect to the variable a uniformly in x and s , and $\mu_e(x, 0, s)$ is uniformly bounded in x and s , and the function ϕ giving the terminal condition is bounded, non-decreasing and Lipschitz, then the FBSDE (30.7) has a unique square integrable solution.

Proof: Let (D_t) and (S_t) represent the strong solutions of the first two equations of (30.7) whose existence is guaranteed by Assumption 30.3.2. These equations being decoupled from the remaining ones, the latter can be treated as a FBSDE with random coefficients and one-dimensional forward and backward components. We claim that existence and uniqueness hold because of Theorem 7.1 of Ma *et al.* (2011).^{*} Strictly speaking, we cannot apply directly Theorem 7.1 of Ma *et al.* (2011) because our Wiener

^{*} We thank Francois Delarue for suggesting this approach.

process is $(n + 1)$ -dimensional. However, a close look at the proof shows that what is really needed is to prove the well-posedness of what the authors call the characteristic BSDE, and the boundedness of its solution and the solutions of the dominating Ordinary Differential Equations (ODE). In the present situation, these equations are rather simple due to the fact that (E_t) has bounded variation, and, as a consequence, its volatility vanishes. The two dominating ODEs can be solved explicitly and one can check that the solutions are bounded by inspection. Moreover, the function ϕ giving the terminal condition being uniformly Lipschitz, the characteristic BSDE is one-dimensional, and although driven by a multi-dimensional Brownian motion, its terminal condition is bounded, and Kobylanski's comparison results (see the original contribution (Kobylanski 2000)) can be used to conclude the proof. ■

The above result is proven for a terminal condition given by a smooth function ϕ . However, as already mentioned earlier, competitive equilibrium arguments for single compliance period models suggest that the function ϕ should be singular (see, for example, Carmona *et al.* (2010)). Indeed, in the event of non-compliance, that is, when the cumulative emissions strictly exceed the cap at the end of the compliance period, i.e. when $E_T > \Gamma$, the penalty and the allowance certificate are perfect substitutes; therefore, they ought to have the same monetary value and one should have $A_T = \pi$, which suggests $\phi(e) = \pi$ whenever $e > \Gamma$. Similarly, in the event of compliance, that is, when the cumulative emissions are strictly below the cap, there will be spare certificates in the market; these certificates will be in zero demand and will therefore expire worthless, so, in this case, $A_T = 0$, which suggests $\phi(e) = 0$ whenever $e < \Gamma$. This economic interpretation of the function ϕ giving the terminal condition gives the whole story when the event $\{E_T = \Gamma\}$ has zero probability since we do not have to worry about the definition of $\phi(e)$ when $e = \Gamma$. Hence the importance of knowing if the random variable E_T is continuous (e.g. has a density). Again, see an early discussion of this property in Carmona *et al.* (2010), and a systematic analysis in Carmona *et al.* (2012b) and Carmona and Delarue (2012). We conjecture that a proof in the spirit of the one given by Carmona and Delarue (2012) should work in the setting of this paper if μ_e is strictly decreasing in a , providing existence and uniqueness of a solution of the FBSDE when the binary terminal condition is weakened. Furthermore, Carmona and Delarue also proved that, still under strict monotonicity of μ_e , the aggregate emissions are equal to the cap with positive probability at the end of the compliance period. This shows that the competitive equilibrium argument given earlier is enough to specify a unique emission process (E_t) and a unique price process (A_t) for the allowance, even though the terminal price of an allowance A_T at the end of the compliance period cannot be prescribed *ex ante* on a set of scenarios of positive probability. We suspect that this is also the case in the present situation.

Note added in proof: The conjectured existence and uniqueness of a solution to the FBSDE (30.7) in our setting was recently proved by Schwarz (2012).

30.4 Valuing Clean Spread Options

In this section we consider the problem of spread option pricing as described in the introduction. Whether the goal is to value a physical asset or to manage the risk associated with financial positions, one needs to compute the price of a European call option on the difference between the price of electricity and the costs of production for a particular power plant. The costs that we take into account are the fixed operation and maintenance costs, the cost of the fuel needed to generate one MWh of electricity and the cost of the ensuing emissions. Letting the $\mathcal{F}_t$ -adapted process (P_t) denote the spot price of electricity, and recasting the informal discussion in the introduction with the notation we chose to allow for several input fuels, a clean spread option with maturity $\tau \in [0, T]$ is characterized by the terminal pay-off

$$(P_\tau - h_i S_\tau^i - e_i A_\tau - K)^+,$$

where K represents the value of the fixed operation and maintenance costs and, for $i \in 1, \dots, n$, $h_i \in \mathbb{R}_{++}$ and $e_i \in \mathbb{R}_{++}$ denote the specific heat and emission rates of the power plant under consideration, and S^i is

the price at time τ of the fuel used in the production of electricity. In the special case when S^i is the price of coal (gas) the option is known as a *clean dark (spark) spread* option.

Since we are pricing by expectation, for $i \in \{1, \dots, n\}$, the value V_t^i of the clean spread is given by the conditional expectation under the pricing measure of the discounted pay-off, i.e.

$$V_t^i = \exp(-r(\tau-t)) \mathbb{E}^Q[(P_\tau - h_i S_\tau^i - e_j A_\tau - K)^+ | \mathcal{F}_t],$$

for $t \in [0, \tau]$.

30.5 A Concrete Two-Fuel Model

We now turn to the special case of two fuels, coal and gas, which differ significantly in their level of emissions per MWh of power generated.

30.5.1 The Bid Stack

Our bid stack model is a slight variation of the one we proposed in Carmona *et al.* (2012a). Here we extend it to include the cost of emissions as part of the variable costs driving firms' bids.

We assume that the coal and gas generators have aggregate capacities $\bar{x}_c$ and $\bar{x}_g$, respectively, so that the market capacity is $\bar{x} = \bar{x}_c + \bar{x}_g$, and their bid levels are given by linear functions of the allowance price and the price of the fuel used for the generation of electricity. We denote these bid functions by b_c and b_g , respectively. The coefficients appearing in these linear functions correspond to the marginal emission rate (measured in ton equivalent of CO₂ per MWh) and the heat rate (measured in MMBtu per MWh) of the technology in question. Specifically, for $i \in \{c, g\}$, we assume that

$$b_i(x, a, s) := e_i(x)a + h_i(x)s, \quad \text{for } (x, a, s) \in [0, \bar{x}_i] \times \mathbb{R} \times \mathbb{R}, \quad (30.8)$$

where the *marginal emission rate* e_i and the *heat rate* h_i are given by

$$\begin{aligned} e_i(x) &:= \hat{e}_i \exp(m_i x), \\ h_i(x) &:= \hat{h}_i \exp(m_i x), \end{aligned} \quad \text{for } x \in [0, \bar{x}_i].$$

Here $\hat{e}_i$, $\hat{h}_i$ and m_i are strictly positive constants. We allow the marginal emission rate and the heat rate of each technology to vary to reflect differences in efficiencies within the fleet of coal and gas generators. Less efficient plants with higher heat rates have correspondingly higher emission rates. We assume that, for each technology, the ratio h_i/e_i is fixed, a reasonable approximation which implies that the emissions rate of any coal (gas) plant is simply a fixed multiple of the quantity of coal (gas) burned.

Proposition 30.5.1: *With b_c and b_g as above and $I = \{c, g\}$, the market bid stack b is given by*

$$b(x, a, s) = \begin{cases} (\hat{e}_i a + \hat{h}_i s_i) \exp(m_i x), & \text{if } b_i(x, a, s_i) \leq b_j(0, a, s_j) \\ & \text{for } i, j \in I, i \neq j, \\ (\hat{e}_i a + \hat{h}_i s_i) \exp(m_i (x - \bar{x}_j)), & \text{if } b_i(x - \bar{x}_j, a, s_i) > b_j(0, a, s_j) \\ & \text{for } i, j \in I, i \neq j \\ \prod_{i \in I} (\hat{e}_i a + \hat{h}_i s_i)^{\beta_i} \exp(\gamma x), & \text{otherwise,} \end{cases}$$

for $(x, a, s) \in [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^2$, where $\beta_i = m_{I \setminus \{i\}} / (m_c + m_g)$ and $\gamma = m_c m_g / (m_c + m_g)$.

Proof: The proof is a straightforward extension of corollary 1 of Carmona *et al.* (2012a).

30.5.2 The Emission Stack

In order to determine the rate at which the market emits we need to know which generators are supplying electricity at any time. By the merit order assumption the market operator calls upon firms in increasing order of their bid levels. Therefore, given electricity, allowance and fuel prices $(p, a, s) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^2$, for $i \in \{c, g\}$, the set of active generators of fuel type i is in one-to-one correspondence with the set $\{x \in [0, \bar{x}_i] : b_i(x, a, s) \leq p\}$.

Proposition 30.5.2: *Assuming that the market bid stack is of the form specified in proposition 5.1, the market emission rate μ_e is given by*

$$\mu_e(x, a, s) := \sum_{i \in \{c, g\}} \frac{\hat{e}_i}{m_i} (\exp(m_i \hat{b}_i^{-1}(b(x, a, s), a, s_i)) - 1), \quad (30.9)$$

for $(x, a, s) \in [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^2$, where, for $i \in \{c, g\}$, we define

$$\hat{b}_i^{-1}(p, a, s_i) := 0 \vee \left(\bar{x}_i \wedge \frac{1}{m_i} \log \left(\frac{p}{\hat{e}_i a + \hat{h}_i s_i} \right) \right),$$

for $(p, a, s) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^2$, and, as usual, $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$.

Proof: The market emission rate follows from integrating the marginal emission rate e_i for each technology over the corresponding set of active generators and then summing the two. Given the monotonicity of b_i in x and its range $[0, \bar{x}_i]$, the function $\hat{b}_i^{-1}$ describes the quantity of electricity supplied by fuel $i \in \{c, g\}$, and hence the required upper limit of integration. ■

30.5.3 Specifying the Exogenous Stochastic Factors

The demand process. We posit that, under $\mathbb{Q}$, the process (D_t) satisfies for $t \in [0, T]$ the stochastic differential equation

$$\begin{aligned} dD_t &= -\eta(D_t - \bar{D}(t))dt + \sqrt{2\eta\hat{\sigma}D_t(\bar{x} - D_t)}d\tilde{W}_t, \\ D_0 &= d_0 \in (0, \bar{x}), \end{aligned}$$

where $[0, T] \ni t \mapsto \bar{D}(t) \in (0, \bar{x})$ is a deterministic function giving the level of mean reversion of the demand and $\eta, \hat{\sigma} \in \mathbb{R}_{++}$ are constants. With this definition (D_t) is a Jacobi diffusion process; it has a linear, mean-reverting drift component and degenerates on the boundary. Moreover, subject to $\min(\bar{D}(t), \bar{x} - \bar{D}(t)) \geq \bar{x}\hat{\sigma}$, for $t \in [0, T]$, the process remains within the interval $(0, \bar{x})$ at all times (Forman and Sørensen 2008). To capture the seasonal character of demand, we choose a function $\bar{D}(t)$ of the form

$$\bar{D}(t) := \varphi_0 + \varphi_1 \sin(2\pi\vartheta t),$$

where the values of the coefficients will be specified in the next section.

The fuel price processes. We assume that the prices of coal (S_t^c) and gas (S_t^g) follow correlated exponential (or geometric) Ornstein–Uhlenbeck processes under the measure $\mathbb{Q}$, i.e. for $i \in \{c, g\}$ and $t \in [0, T]$,

$$\begin{aligned} dS_t^i &= -\eta_i \left(\log S_t^i - \bar{s}_i - \frac{\hat{\sigma}_i^2}{2\eta_i} \right) S_t^i dt + \hat{\sigma}_i S_t^i d\tilde{W}_t^i, \\ S_0^i &= s_0^i \in \mathbb{R}_{++}, \end{aligned}$$

where $d\langle W^c, W^g \rangle_t = \rho dt$.

30.6 Numerical Analysis

We now turn to the detailed analysis of the model we propose. For this purpose we consider a number of case studies in Sections 30.6.2 to 30.6.5. To produce the following results we used the numerical schemes explained in Appendices A and B.

30.6.1 Choice of Parameters

The tables in this section specify the values of the parameters used for the numerical analysis of our model that follows below. We refer to the parameter values specified in Tables 30.1–30.5 as the ‘base case’ and indicate whenever we depart from this choice. Note that our choices do not correspond to a particular electricity market, but that all values are within a realistic realm.

Table 30.1 provides the parameter values specifying the bid curves. We consider a medium-sized electricity market served by coal and gas generators and with gas being the dominant technology. For the marginal emission rates, Table 30.1 implies that $e_c \in [0.9, 1.64]$ and $e_g \in [0.4, 0.69]$ (both measured in tCO₂ per MWh), so that all gas plants are ‘cleaner’ than all coal plants. For the heat rates, we observe that $h_c \in [3, 5.5]$ and $h_g \in [7, 12]$ (both measured in MMBtu per MWh). Using (30.9) now with $D_t = \bar{x}$, for $0 \leq t \leq T$, and the assumption that there are 8760 production hours in the year, we find, denoting the maximum cumulative emissions by $\bar{e}$, that $\bar{e} = 2.13 \times 10^8$.

Table 30.2 contains the parameter values for the demand process (D_t). We model periodicities on an annual and a weekly time scale and the chosen rate of mean-reversion assumes that demand reverts to its (time-dependent) mean over the course of one week.

In Table 30.3 we give the parameter values that specify the behaviour of the prices of coal and gas. Both are chosen to be slowly mean-reverting, at least in comparison to demand. To ease the analysis, we assume that all parameters are identical for the two fuels, including mean price levels, both measured in MMBtu.*

TABLE 30.1 Parameters Relating to the Bid and Emission Stacks

$\hat{h}_c$	$\hat{e}_c$	m_c	$\bar{x}_c$	$\hat{h}_g$	$\hat{e}_g$	m_g	$\hat{x}_g$	$\bar{x}$
3	0.9	0.00005	12,000	7	0.4	0.00003	18,000	30,000

TABLE 30.2 Parameters Relating to the Demand Process

η	φ_0	φ_1	ϑ	$\hat{\sigma}$	d_0
50	21,000	3000	1	0.1	21,000

TABLE 30.3 Parameters Relating to the Fuel Price Processes

η_c	$\bar{s}_c$	$\hat{\sigma}_c$	s_0^c	η_g	$\bar{s}_g$	$\hat{\sigma}_g$	s_0^g	ρ
1.5	2	0.5	exp(2)	1.5	2	0.5	exp(2)	0.3

* We note that gas and coal prices are typically quoted in different units, and can often differ by a factor of 10 or more. However, in our analysis, as we are not fitting to data, coal and gas only play the role of common representative fuel types (and other possibilities include lignite, oil, etc.). Therefore, our parameter choices simply reflect typical characteristics of energy price behavior. Much more importantly for our analysis, we require that one fuel be significantly ‘cleaner’ than the other, and that the relative price levels allow for merit order changes driven by the cap-and-trade market.

TABLE 30.4 Parameters Relating to the Cap-and-Trade Scheme

π	Γ	T	r
100	1.4×10^8	1	0.05

TABLE 30.5 Parameters Relating to the Spread Options

High eff. coal		Low eff. coal		High eff. gas		Low eff. gas	
h_c	e_c	h_c	e_c	h_g	e_g	h_g	e_g
3.5	1.05	5.0	1.5	7.5	0.43	11.5	0.66

Table 30.4 defines the cap-and-trade scheme that we assume to be in place. The duration of the compliance period T is measured in years and we set the cap at 70% of the upper bound $\bar{e}$ for the cumulative emissions, in order to incentivise a reduction in emissions. This choice of parameter values results in A_0 being approximately equal to $\pi/2$, a value for which there is significant initial overlap between gas and coal bids in the stack. Furthermore, the values imply a bid stack structure such that, at mean levels of coal and gas prices, $A_t = 0$ pushes all coal bids below gas bids, while for $A_t = \pi$ almost all coal bids are above all gas bids.

Finally, in Table 30.5 we specify the four spread option contracts used in the base case scenario to represent high and low efficiency coal plants, and high and low efficiency gas plants (note that low (high) efficiency means dirtier (cleaner) and corresponds to high (low) h_i and e_i).

We now consider a series of case studies to investigate various features of the model's results in turn. As the model captures many different factors and effects, this allows us to isolate some of the most important implications. In Case study I, we investigate the impact on coal and gas plants of different efficiencies of creating an increasingly strict carbon emissions market. In Case study II, we assess the impact on these plants of changes in initial fuel prices. In Case study III, we compare spread option prices in our model with two simple reduced-form approaches for A_t , which allows us to better understand the role of key model features such as bid-stack-driven abatement. Finally, in case study IV, we consider the overall impact of cap-and-trade markets in the electricity sector, by comparing them with a well-known alternative, a fixed carbon tax.

30.6.2 Case Study I: Impact of the Emission Market

The first effect that we are interested in studying is the impact of the cap-and-trade market on clean spread option prices, for increasingly strict levels of the cap Γ . At one extreme (when the cap is so generous that $A_t \approx 0$, for all $t \in [0, T]$), the results correspond to the case of a market without a cap-and-trade system, while at the other extreme (when the cap is so strict that $A_t \approx \pi \exp(-r(T-t))$, for all $t \in [0, T]$), there is essentially a very high carbon tax which tends to push most coal generators above gas generators in the stack. It is intuitively clear that higher carbon prices typically lead to higher spark spread option prices and lower dark spread option prices, thus favouring gas plants over coal plants, but the relationships can be more involved as we vary between low and high efficiency plants.

In Figure 30.1, we compare spread option prices corresponding to different efficiency generators (i.e. to different h_i, e_i in the spread pay-off) as a function of maturity τ . 'High' and 'low' efficiency plant indicates values of h_i, e_i chosen to be near the lowest and highest, respectively, in the stack, as given by Table 30.5. Within each of the four subplots, the five lines correspond to five different values of the cap Γ , ranging from very lenient to very strict. We immediately observe from Figure 30.1 the seasonality in spread prices caused by the seasonality in power demand. This is most striking for the low efficiency cases (high h_i, e_i), as such plants would rarely be used in shoulder months, particularly in the case of gas. For low efficiency plants, the relationship with cap level (and corresponding initial allowance price) is as one would expect: a stricter cap greatly increases the value of gas plants and greatly decreases the value of the dirtier coal plants. This is also true for high efficiency gas plants, although the price difference (in




FIGURE 30.1 Cap strictness analysis for high efficiency coal (top left), low efficiency coal (top right), high efficiency gas (bottom left) and low efficiency gas (bottom right): spark and dark spread option values plotted against maturity, for varying levels of the cap Γ . Note that the five equally spaced cap values from 1.8×10^8 to 1.0×10^8 tons of CO_2 imply initial allowance prices of \$5, \$28, \$52, \$80 and \$94.

percentage terms) for different Γ is less, since these are effectively ‘in-the-money’ options, unlike those discussed above. However, the analysis becomes more complicated for high efficiency coal plants, which tend to be chosen to run in most market conditions, irrespective of emission markets. Interestingly, we find that, for these options, the relationship with Γ (and hence A_0) can be non-monotonic under certain conditions, particularly for high levels of demand, when the price is set near the very top of the stack. In such cases a stricter cap provides extra benefit for the cleaner coal plants via higher power prices (typically set by the dirtier coal plants on the margin) which outweighs the disadvantage of coal plants being replaced by gas plants in the merit order.

30.6.3 Case Study II: Impact of Fuel Price Changes

Note that, in Table 30.3, the initial conditions of both gas and coal have been set to be equal to their long-term median levels. We now consider the case of the gas price s_0^g being either above or below its long-term level, thus inducing a change in the initial merit order. Given the record low prices of under \$2 recently witnessed in the US natural gas market (due primarily to shale gas discoveries), it is natural to ask how such fuel price variations affect our spread option results. Note, however, that since $\eta_c = \eta_g = 1.5$ (implying a typical mean-reversion time of 8 months), by the end of the trading period the simulated fuel price distributions will again be centred near their mean-reversion levels. Thus in this case study, we capture a temporary, not permanent, shift in fuel prices.

In Figure 30.2, we plot the value of coal and gas power plants, as given by the sum of spread options of all maturities $\tau \in [0, T]$. In the first plot, we consider low efficiency (high h_i and e_i) plants, while in the second we consider high efficiencies. The latter are much more likely to operate each day and to generate profits, and are hence much more valuable than the former. However, they also show different relationships with s_0^g , as illustrated for several different cap levels Γ (like in Case study I above) which correspond




FIGURE 30.2 Power Plant Value (sum of spreads over τ) versus s_0^g for low efficiency (left) and high efficiency (right). 'High A' corresponds to $\Gamma = 1 \times 10^8$, 'mid A' to $\Gamma = 1.4 \times 10^8$ (base case) and 'low A' to $\Gamma = 1.8 \times 10^8$, with corresponding values $A_0 = 94$, $A_0 = 52$ and $A_0 = 5$.

to high, low or medium (base case) values of A_0 .^{*} Firstly, for low efficiency plants (left plot), we observe that the gas plant value is typically decreasing in s_0^g , as we expect, since higher gas prices tend to push the bids from gas above those from coal, meaning there is less chance that the gas plant will be used for electricity generation. Similarly, coal plant values are typically increasing in s_0^g , as more coal plants will be used. Note however that, for some cases, the curves flatten out, as no more merit order changes are possible. This is particularly true for the coal plant when A_0 is very high (and hence once gas drops below a certain point, the coal plant is almost certainly going to remain more expensive to run than all gas plants) and for the gas plant when A_0 is very low (and hence once coal increases above a certain point, the gas plant is almost certainly going to remain more expensive to run than all coal plants).

We now turn our attention to the high efficiency case (right plot), meaning the relatively cheap and clean plants for each technology. As expected, the coal plant benefits from low values of A_0 (as implied by a lenient cap) and gas from high values of A_0 (i.e. a strict cap). On the other hand, the relationship with s_0^g is now increasing for all cases plotted except that of a gas plant with high A_0 . While it may seem surprising that, for low values of A_0 (or medium though not plotted), the gas plant value increases with s_0^g , this is quite intuitive when one considers that the range of bids from gas generators widens as s_0^g increases, implying that the efficient plants can make a larger profit when the inefficient plants set the power price. Indeed, as demand is quite high on average, and gas is 60% of the market, it is likely that these efficient gas plants will almost always be 'in-the-money' even if coal is lower in the stack. Only in the case where coal is typically above gas and now marginal (i.e. the high A_0 case) is the value of the gas plant decreasing in s_0^g since the plant's profit margins shrink as gas and coal bids converge.

30.6.4 Case Study III: Comparison with Reduced-Form

The second analysis we consider is to compare the results of our structural model for the allowance price, with two other simpler models, both of which belong to the class of 'reduced-form' models. The first of these treats the allowance price itself as a simple Geometric Brownian Motion (with drift r under $\mathbb{Q}$), and hence A_t is log-normal at spread maturity, like S_t^c and S_t^g . The second comparison treats the emission process as a Geometric Brownian Motion (GBM), and retains the digital terminal condition $A_T = \pi_{\{E_T \geq \Gamma\}}^0$. As the drift of (E_t) is then simply a constant (chosen to match the initial value A_0 in the full model), there is no feedback from (A_t) to (E_t) , or in other words, no abatement induced by the allowance price. For any time t , A_t is then given in closed form by a formula resembling the Black-Scholes digital option price. In

^{*} In the first plot, the cases 'coal - low A' and 'gas - high A' would produce values much higher than the other cases, and hence we instead choose 'coal - mid A' and 'gas - mid A' in order to illustrate the effects on a single plot.

order to fully specify the two reduced-form models, we need to choose a volatility parameter σ_a or σ_e for each of the GBMs, as well as correlations ρ_{ac} , ρ_{ag} or ρ_{ec} , ρ_{eg} with the Brownian Motions driving the other exogenous factors, coal and gas prices. All of these parameters are chosen to approximately match the levels of volatility and correlation produced by simulations in the full structural model, and are given in Table 30.6. Finally, note that, in all three models that we compare, the power price is given by the same bid stack function as usual, so our aim is to isolate and evaluate the effect of our more sophisticated framework for the allowance price, in comparison with simpler approaches. The cap throughout is $\Gamma = 1.4 \times 10^8$, the base case.

Figure 30.3 reveals that the difference between the reduced-form models and the full structural model is relatively small for high efficiency gas and coal plants which are typically ‘in-the-money’. In contrast, a larger gap appears for low efficiency cases, where the reduced-form models significantly overprice spread options relative to the stack model. In particular, the case of log-normal emissions produces much higher prices, especially for dark spreads. The intuition is as follows. In the full model, the bid stack structure automatically leads to lower emissions when the allowance price is high, and higher emissions when the allowance price is low, producing a mean-reversion-like effect on the cumulative emissions, keeping the process moving roughly towards the cap, with the final outcome (compliance or not) in many simulations only becoming clear very close to maturity. In contrast, if (E_t) is a GBM, much of the uncertainty is often resolved early in the trading period, with (A_t) then sticking near zero or π for much of the period. In such cases, there is a much larger benefit for deep OTM options (low efficiency plants), for which the tails of

TABLE 30.6 Parameters for Reduced-Form Comparisons, Treating A_t and E_t as GBMs

σ_a	ρ_{ac}	ρ_{ag}	σ_e	ρ_{ec}	ρ_{eg}
0.6	-0.2	0.4	0.006	-0.2	0.2




FIGURE 30.3 Model comparison against reduced-form: spark and dark spread option values for varying heat rates, emission rates and maturities.

the allowance price distribution provide great value either for coal (when the price is near zero) or for gas (when the price is near the penalty). We observe that, in some of the subplots (particularly low efficiency coal), this extra benefit is indeed realized in the full model, but only very near the end of the trading period when the volatility of (A_t) often spikes, and the process either rises or falls sharply. In contrast, for the other reduced-form model with log-normally distributed allowance price, the volatility of the allowance price is constant throughout and (A_t) never moves rapidly towards zero or the penalty. However, the overall link with fuel and power prices is much weaker when simply using correlated Brownian Motions, which serves to widen the spread distribution in most cases relative to the full structural model. This result is somewhat similar to the observation of Carmona *et al.* (2012a) that a stack model generally produces lower spread option prices than Margrabe's formula for correlated log-normals.

30.6.5 Case Study IV: Cap-and-Trade vs. Carbon Tax

Finally, we wish to investigate the implications of the model for cap-and-trade systems, as compared with fixed carbon taxes. This question has been much debated by policy makers as well as academics, and can be roughly summarized as fixing quantity versus fixing price. Carmona *et al.* (2010) compare several different designs for cap-and-trade systems with a carbon tax, using criteria such as cost to society and windfall profits to power generators. Here we follow a related approach by analysing the power sector as a whole, but we build on our previous case studies by using spread option prices as a starting point. Firstly, we observe that the total expected discounted profits of the power sector are equal to the value of all the power plants implied by the bid stack structure, which in turn equals a portfolio of (or integral over) sums of spread option prices with varying h_i and e_i . Therefore, for each simulation over the period $[0, T]$, total profits (total revenues minus total costs) are*

$$\begin{aligned} \text{Total profits} &= \sum_{\tau \in [0, T]} \left(P_\tau D_\tau - \int_0^{D_\tau} b(x, A_\tau, S_\tau) dx \right) \\ &= \sum_{\tau \in [0, T]} \int_0^{\bar{x}} P_\tau - b(x, A_\tau, S_\tau) dx \\ &= \sum_{\tau \in [0, T]} \left(\int_0^{\bar{x}_c} P_\tau - h_c(x) S_\tau^c - e_c(x) A_\tau dx \right. \\ &\quad \left. + \int_0^{\bar{x}_g} (P_\tau - h_g(x) S_\tau^g - e_g(x) A_\tau) dx \right), \end{aligned}$$

where the second line follows from the fact that the events $\{P_\tau \geq b(x, A_\tau, S_\tau)\}$ and $\{D_\tau \geq x\}$ are equal.

Hence, instead of picking particular coal and gas plants with efficiencies specified by the parameters in Table 30.5, we now integrate power plant value over all the efficiencies of plants in the stack, as defined by the parameters in Table 30.1. In the case of the carbon tax, we simply force $A_t = A_0 \exp(rt)$ for all $t \in [0, T]$, including the exponential function in order to match the mean of the process in the cap-and-trade model. This is equivalent to setting the volatility σ_a equal to zero in the GBM model for the allowance price in Case study III.

In Figure 30.4, we first plot the expected total market profits in the base case as a function of time. It is interesting to observe that two important effects occur, pulling the profits in opposite directions, but varying in strength over the trading period. In particular, although the profits must be equal at time zero, a gap quickly appears in the early part of the trading period, with expected profits to power generators significantly higher under a carbon tax than cap-and-trade. However, as maturity approaches, the gap narrows and the order reverses over the final days, as cap-and-trade generates higher expected profits. We reason as follows: firstly, as $A_0 = 52$ in the base case, the bids of coal and gas begin the period

* Note that we do not consider here additional issues such as whether allowances are auctioned or freely allocated to generators. Instead, we assume that allowances are bought on the market by generators as and when they need them.




FIGURE 30.4 Cap-and-trade vs. carbon tax: power sector profits versus time for the 'base case' (left); total profits over one year for equally spaced cap values from 1×10^8 to 1.95×10^8 tons of CO₂ (right).

at very similar levels, a state which generally keeps profits low, since the variance of electricity prices is low and the profit margins of both coal and gas generators are quite low. As time progresses and fuel prices move, the coal and gas bids will tend to drift apart in most simulations, for example with gas sometimes moving above coal, say. However, in our structural model for the cap-and-trade scheme, in such a case the higher emissions will induce a higher allowance price, and in turn a feedback effect due to the coupling in (30.7), which acts to keep coal and gas bids closer together. A similar argument can be made for the case of gas bids tending to move below coal bids but then being counteracted by lower allowance prices. Again we see that the power market structure induces mean-reversion on (E_t), which in this scenario (of a mid-range cap level) corresponds to keeping coal and gas bids close together. On the other hand, under a carbon tax with fixed (or deterministic) allowance price, there is of course no feedback mechanism (no price-sensitive abatement), and bids tend to wander apart. However, as the end of the trading period approaches, in the cap-and-trade system the allowance price eventually gets pulled to either zero or the penalty, which will separate the bids in one way or the other, either leading to very large profits for coal plants (if $A_T = 0$) or for gas plants (if $A_T = \pi$). This is a similar effect to that discussed when comparing with a log-normal allowance price in Case study III, as neither a carbon tax nor a log-normal allowance price model sees the extra volatility near maturity caused by the terminal condition.

Finally, in the second plot of Figure 30.4, we consider how these conclusions change if the cap is made stricter or more lenient. Instead of plotting against maturity, we consider the total profits of the power sector over the entire period $[0, T]$. Firstly, we observe that, under both forms of emission regulation, power sector profits are lowest if the cap is chosen close to the base case, under which the bids from coal and gas generators are more tightly clustered together. Secondly, it is important to notice that the conclusion in the previous discussion that a carbon tax provides more profits to the power sector does not hold for all scenarios of the cap. In particular, for either very high or very low values of the cap, the cap-and-trade scheme provides more profits than a tax. The explanation here is that for the automatic abatement mechanism in the stack to have its largest impact (keeping bids together, and emissions heading towards the cap), there needs to be significant uncertainty at time zero as to whether the cap will be reached. The feedback mechanism of a cap-and-trade system then allows this uncertainty to be prolonged through the period. On the other hand, for an overly strict or overly lenient cap (or similarly for a merit order which does not allow for much abatement), the second effect discussed above dominates over the first. In other words, the terminal condition which guarantees large profits to either coal or gas at maturity begins to take precedence earlier in the trading period, instead of just before maturity as in the base case. Although in practice there are many other details to consider when comparing different forms of emission legislation, our stylized

single-period model sheds some light on the differences between cap-and-trade and carbon tax, as well as the clear importance of choosing an appropriate cap level.

30.7 Conclusion

As policy makers debate the future of global carbon emission legislation, the existing cap-and-trade schemes around the world have already significantly impacted the dynamics of electricity prices and the valuation of real assets, such as power plants, particularly under the well-known European Union Emissions Trading Scheme. Together with the recent volatile behaviour of all energy prices (e.g. gas, coal, oil), the introduction of carbon markets has increased the risk of changes in the merit order of fuel types, known to be a crucial factor in the price-setting mechanism of electricity markets. In the US, the recent sharp drop in natural gas prices is already causing changes in the merit order, which would be further magnified by any new emission regulation, such as the upcoming cap-and-trade market in California. Such considerations are vital for describing the complex dependence structure between electricity, its input fuels, and emission allowances, and thus highly relevant for both market participants and policy makers designing emission trading schemes. In this paper, we derived the equilibrium carbon allowance price as the solution of an FBSDE, in which feedback from allowance price on market emission rates is linked to the electricity stack structure. The resulting model specifies simultaneously both electricity and allowance price dynamics as a function of fuel prices, demand and accumulated emissions; in this way, it captures consistently the highly state-dependent correlations between all the energy prices, which would not be achievable in a typical reduced-form approach. We used a PDE representation for the solution of the pricing FBSDE and implemented a finite difference scheme to solve for the price of carbon allowances. Finally, we compared our model for allowance prices with other reduced-form approaches and analysed its important implications on price behaviour, spread option pricing and the valuation of physical assets in electricity markets covered by emission regulation. The four case studies illustrate the many important considerations needed to understand the complex joint dynamics of electricity, emissions and fuels, as well as the additional insight that can be provided by our structural approach.

Acknowledgements

R.C. is partially supported by NSF-DMS 0806591. M.C. is partially supported by NSF-DMS-0739195.

References

- Aïd, R., Campi, L. and Langrené, N., A structural risk-neutral model for pricing and hedging power derivatives. *Math. Finance*, 2012, to appear.
- Alos, E., Eydeland, A. and Laurence, P., A Kirks and a Bacheliers formula for three-asset spread options. *Energy Risk*, 2011, **September**, 52–57.
- Barlow, M., A diffusion model for electricity prices. *Math. Finance*, 2002, **12**, 287–298.
- Benth, F.E., Benth, J.S. and Koekebakker, S., *Stochastic Modeling of Electricity and Related Markets*, 2008 (World Scientific: Singapore).
- Carmona, R. and Coulon, M., A survey of commodity markets and structural models for electricity prices. In *Financial Engineering for Energy Asset Management and Hedging in Commodity Markets. Proceedings from the Special Thematic Year at the Wolfgang Pauli Institute, Vienna*, edited by F.E. Benth, 2012 (Springer).
- Carmona, R. and Delarue, F., Singular FBSDEs and scalar conservation laws driven by diffusion processes. *Probab. Theory Relat. Fields*, 2012, to appear.
- Carmona, R. and Durrleman, V., Pricing and hedging spread options. *SIAM Rev.*, 2003, **45**(4), 627–685.
- Carmona, R. and Hinz, J., Risk neutral modeling of emission allowance prices and option valuation. *Mgmt Sci.*, 2011, **57**(8), 1453–1468.

- Carmona, R. and Sun, Y., Implied and local correlations from spread options. *Appl. Math. Finance*, 2012.
- Carmona, R., Coulon, M. and Schwarz, D., Electricity price modelling and asset valuation: A multi-fuel structural approach. *Math. Financ. Econ.*, 2012a, to appear.
- Carmona, R., Delarue, F., Espinoza, G.E. and Touzi, N., Singular forward backward stochastic differential equations and emission derivatives. *Ann. Appl. Probab.*, 2012b.
- Carmona, R., Fehr, F., Hinz, J. and Porchet, A., Market designs for emissions trading schemes. *SIAM Rev.*, 2010, **52**, 403–452.
- Cartea, Á. and Villaplana, P., Spot price modeling and the valuation of electricity forward contracts: The role of demand and capacity. *J. Bank. Finance*, 2008, **32**, 2501–2519.
- Chesney, M. and Taschini, L., The endogenous price dynamics of the emission allowances: An application to CO₂ option pricing. *Appl. Math. Finance*, 2012, forthcoming.
- Coulon, M., Modelling price dynamics through fundamental relationships in electricity and other energy markets. PhD thesis, University of Oxford, 2009.
- Coulon, M. and Howison, S., Stochastic behaviour of the electricity bid stack: From fundamental drivers to power prices. *J. Energy Mkts*, 2009, **2**, 29–69.
- De Jong, C. and Schneider, S., Cointegration between gas and power spot prices. *J. Energy Mkts*, 2009, **2**(3), 27–46.
- Eydeland, A. and Wolyniec, K., *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*, 2003 (Wiley: New York).
- Forman, J.L. and Sørensen, M., The Pearson diffusions: A class of statistically tractable diffusion processes. *Scand. J. Statist.*, 2008, **35**, 438–465.
- Glasserman, P., *Monte Carlo Methods in Financial Engineering*, 2004 (Springer: Berlin).
- Howison, S. and Schwarz, D., Risk-neutral pricing of financial instruments in emission markets: A structural approach. *SIAM J. Financ. Math.*, 2012, to appear.
- Karatzas, I. and Shreve, S.E., *Brownian Motion and Stochastic Calculus*, 1999 (Springer: Berlin).
- Kobylanski, M., Backward stochastic differential equations and partial differential equations with quadratic growth. *Ann. Probab.*, 2000, **28**, 558–602.
- Koenig, P., Modelling correlation in carbon and energy markets. Working Paper, Department of Economics, University of Cambridge, 2011.
- LeVeque, R.J., *Numerical Methods for Conservation Laws*, 1990 (Birkhäuser).
- Ma, J., Wu, Z., Zhang, D. and Zhang, J., On well-posedness of forward-backward SDEs|A unified approach. Working Paper, University of Southern California, 2011.
- Margrabe, W., The value of an option to exchange one asset for another. *J. Finance*, 1978, **33**, 177–186.
- Oleinik, O.A. and Radkevich, E.V., *Second Order Equations with Nonnegative Characteristic Form*, 1973 (AMS).
- Pirrong, C. and Jermakyan, M., The price of power: The valuation of power and weather derivatives. *J. Bank. Finance*, 2008, **32**, 2520–2529.
- Schwarz, D., Price modelling and asset valuation in carbon emission and electricity markets. PhD thesis, University of Oxford, 2012.
- Seifert, J., Uhrig-Homburg, M. and Wagner, M.W., Dynamic behaviour of CO₂ spot prices. *J. Environ. Econ. Mgmt*, 2008, **56**(2), 180–194.

Appendix A: Numerical Solution of the FBSDE

A.1 Candidate Pricing PDE

In Theorem 30.3.3 we addressed the existence and uniqueness of a solution to the FBSDE (30.7). Given the Markov nature of the equation, we conjecture that there exists a deterministic function $\alpha: [0, T] \times [0, \bar{x}] \times \mathbb{R}_{++} \times [0, \bar{e}] \hookrightarrow [0, \pi]$ such that $A_t = \alpha(t, D_t, E_t, S_t^e, S_t^g)$, and sufficiently smooth to be a classical solution to the semilinear PDE

$$\mathcal{L}\alpha + \mathcal{N}\alpha = 0, \quad \text{on } U_T, \quad (\text{A.1})$$

$$\alpha = \phi(e), \quad \text{on } \{t = T\} \times U, \quad (\text{A.2})$$

where $U := (0, \bar{x}) \times \mathbb{R}_{++} \times \mathbb{R}_{++} \times (0, \bar{e})$ and $U_T := [0, T] \times U$ and the operators $\mathcal{L}$ and $\mathcal{N}$ are defined by

$$\begin{aligned} \mathcal{L} := & \frac{\partial}{\partial t} + \frac{1}{2} \sigma_d(d)^2 \frac{\partial^2}{\partial d^2} + \frac{1}{2} \sigma_c(s_c)^2 \frac{\partial^2}{\partial s_c^2} + \frac{1}{2} \sigma_g(s_g)^2 \frac{\partial^2}{\partial s_g^2} \\ & + \rho \sigma_c(s_c) \sigma_g(s_g) \frac{\partial^2}{\partial s_c \partial s_g} \\ & + \mu_d(t, d) \frac{\partial}{\partial d} + \mu_c(s_c) \frac{\partial}{\partial s_c} + \mu_g(s_g) \frac{\partial}{\partial s_g} - r, \end{aligned}$$

and $\mathcal{N} := \mu_e(d, \cdot, (s_c, s_g))(\partial/\partial e)$. As previously, we specify for our purposes that $\phi(e) = \pi I_{[\Gamma, \infty)}(e)$, for $e \in \mathbb{R}$.

With regards to the problem (A.1), the question arises at which parts of the boundary we need to specify boundary conditions and, given the original stochastic problem (30.7), of what form these conditions should be. To answer the former question we consider the Fichera function f at points of the boundary where one or more of the diffusion coefficients disappear (Oleinik and Radkevich 1973). Defining $n := (n_d, n_c, n_g, n_e)$ to be the inward normal vector to the boundary, Fichera's function for the operator $(\mathcal{N} + \mathcal{L})$ reads

$$\begin{aligned} f(t, d, s_c, s_g, e) := & \left(\mu_d - \frac{1}{2} \frac{\partial}{\partial d} \sigma_d^2 \right) n_d \\ & + \left(\mu_c - \frac{1}{2} \frac{\partial}{\partial s_c} \sigma_c^2 - \frac{\partial}{\partial s_c} \rho \sigma_c \sigma_g \right) n_c \\ & + \left(\mu_g - \frac{1}{2} \frac{\partial}{\partial s_g} \sigma_g^2 - \frac{\partial}{\partial s_g} \rho \sigma_c \sigma_g \right) n_g \\ & + \mu_e n_e, \quad \text{on } \partial U_T. \end{aligned} \quad (\text{A.3})$$

At points of the boundary where $f \geq 0$ the direction of information propagation is outward and we do not need to specify any boundary conditions; at points where $f < 0$ information is inward flowing and boundary conditions have to be specified. We evaluate (A.3) for the choice of coefficients presented in Section 30.5.3.

Considering the parts of the boundary corresponding to $d = 0$ and $d = \bar{x}$, we find that $f \geq 0$ if and only if $\min(\bar{D}(t), \bar{x} - \bar{D}(t)) \geq \bar{x} \hat{\sigma}$, which is the same condition prescribed in Section 30.5.3 to guarantee that the Jacobi diffusion stays within the interval $(0, \bar{x})$. At points of the boundary corresponding to $e = 0$, we find that $f \geq 0$ always. On the part of the boundary on which $e = \bar{e}$, $f < 0$ except at the point $(d, \cdot, \cdot, e) = (0, \cdot, \cdot, \bar{e})$, where $f = 0$, an ambiguity which could be resolved by smoothing the domain. Similarly, we find that $f \geq 0$ on parts of the boundary where $s_c = 0$ or $s_g = 0$. Therefore, no boundary conditions are necessary except when $e = \bar{e}$, where we prescribe

$$\alpha = \exp(-r(T-t))\pi, \quad \text{on } U_T|_{e=\bar{e}}. \quad (\text{A.4})$$

In addition, we need to specify an asymptotic condition for large values of s_c and s_g . We choose to consider solutions that, for $i \in \{c, g\}$, satisfy

$$\frac{\partial \alpha}{\partial s_i} \sim 0, \quad \text{on } U_T|_{s_i \rightarrow \infty}. \quad (\text{A.5})$$

A.2 An Implicit–Explicit Finite Difference Scheme

We approximate the domain U_T by a finite grid spanning $[0, T] \times [0, \bar{x}] \times [0, \bar{s}_c] \times [0, \bar{s}_g] \times [0, \bar{e}]$. For the discretization we choose mesh widths Δd , Δs_c , Δs_g , Δe and a time step Δt . The discrete mesh points $(t_k, d_m, s_{c_i}, s_{g_j}, e_n)$ are then defined by

$$\begin{aligned} t_k &:= k\Delta t, & d_m &:= m\Delta d, \\ s_{c_i} &:= i\Delta s_c, & s_{g_j} &:= j\Delta s_g, & e_n &:= n\Delta e. \end{aligned}$$

The finite difference scheme we employ produces approximations $\alpha_{m,i,j,n}^k$ which are assumed to converge to the true solution α as the mesh width tends to zero.

Since the partial differential equation (A.1) is posed backwards in time with a terminal condition, we choose a backward finite difference for the time derivative. In order to achieve better stability properties we make the part of the scheme relating to the linear operator $\mathcal{L}$ implicit; the part relating to the operator $\mathcal{N}$ is made explicit in order to handle the nonlinearity.

In the e -direction we are approximating a conservation law PDE with discontinuous terminal condition (for an in-depth discussion of numerical schemes for this type of equation, see LeVeque (1990)). The first derivative in the s -direction, relating to the nonlinear part of the partial differential equation, is discretized against the drift direction using a one-sided upwind difference. Because characteristic information is propagating in the direction of decreasing e , this one-sided difference is also used to calculate the value of the approximation on the part of the boundary corresponding to $e = 0$. On the part of the boundary corresponding to $e = \bar{e}$ we apply the condition (A.4).

In the d -direction the equation is elliptic everywhere except on the boundary, where it degenerates. Therefore, we expect the convection coefficient to be much larger than the diffusion coefficient near the boundaries. In order to keep the discrete maximum principle we again use a one-sided upwind difference for the first-order derivative. Thereby we have to pay attention that, due to the mean-reverting nature of (D_t) , the direction of information propagation and therefore the upwind direction changes as the sign of μ_d changes. The same upwind difference is also used to calculate the value of the approximation at the boundaries $d = 0$ and $d = \bar{x}$. To discretize the second-order derivative we use central differences.

The s_c - and s_g -directions are treated similarly to the d -direction. We use one-sided upwind differences for the first-order derivatives, thereby taking care of the boundaries corresponding to $s_c = 0$ and $s_g = 0$. The second-order derivatives are discretized using central differences. At the boundary corresponding to $s_c = \bar{s}_c$ and $s_g = \bar{s}_g$ we apply the asymptotic condition (A.5) as a boundary condition.

With smooth boundary data, on a smooth domain, the scheme described above can be expected to exhibit first-order convergence. In our setting, we expect the discontinuous terminal condition to have adverse effects on the convergence rate. We refer the interested reader to Howison and Schwarz (2012) for a numerical estimation and a detailed analysis of the convergence rate of the numerical scheme described in this section, as applied to a very similar case.*

Appendix B: Numerical Calculation of Spread Prices

B.1 Time Discretization of SDEs

Let $(D_k, S_k^c, S_k^g, E_k, A_k)$ denote the discrete-time approximation to the FBSDE solution $(D_t, S_t^c, S_t^g, E_t, A_t)$ on the time grid $0 < \Delta t < 2\Delta t < \dots < n_k \Delta t = \tau$. At each time step we calculate A_k by interpolating the discrete approximation $\alpha_{m,i,j,n}^k$ at (D_k, S_k^c, S_k^g, E_k) beginning with the initial values $D_0 = d_0$, $S_0^c = s_0^c$, $S_0^g = s_0^g$,

* Although the setting of Howison and Schwarz (2012) leads to a 2 + 1-dimensional model instead of the 4 + 1-dimensional model we are considering here, we emphasize that structurally the two models are identical. Therefore, we believe that the convergence properties of the two numerical schemes are also similar.

$E_0 = 0$. The approximations (D_k, S_k^c, S_k^g, E_k) are obtained using a simple Euler scheme (Glasserman 2004) for the forward components of (30.7). The discretized version of (D_t) is forced to be instantaneously reflecting at the boundaries $D_k = 0$ and $D_k = \bar{x}$; similarly, the discretized versions of (S_t^c) and (S_t^g) are made instantaneously reflecting at $S_k^c = 0$ and $S_k^g = 0$.

B.2 Monte Carlo Calculation of Option Prices

Using this discretization we simulate n_{mc} paths and, as usual, for $t \in [0, \tau)$, calculate the mean clean spread price $\hat{V}_t^j$, given by

$$\hat{V}_t^j := \exp(-r(\tau-t)) \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (b(D_{n_k}^i, S_{n_k}^{c,i}, S_{n_k}^{g,i}, A_{n_k}^i) - h_j S_{n_k}^{j,i} - e_j A_{n_k}^i)^+,$$

where the index i refers to the simulation scenario and $j \in \{c, g\}$. The corresponding standard error $\hat{\sigma}_v$ is obtained by

$$\hat{\sigma}_v := \sqrt{\frac{1}{n_{mc}(n_{mc}-1)} \sum_{i=1}^{n_{mc}} (V_{n_k}^i - \hat{V}_\tau^j)^2}.$$

Is the EUA a New Asset Class?

Vicente Medina

Angel Pardo

31.1	Introduction	667
31.2	Stylised Facts of Asset Returns	668
31.3	Data	669
31.4	Results.....	671
	Data Distribution • Volatility-Related Features • Commodity Features	
31.5	Conclusions.....	685
	Acknowledgements	685
	References.....	686

The listing of a new asset requires knowledge of its statistical properties prior to its use for hedging, speculative or risk management purposes. In this paper, the authors study the stylised facts of European Union Allowances (EUAs) returns. The majority of the phenomena observed, such as heavy tails, volatility clustering, asymmetric volatility and the presence of a high number of outliers are similar to those observed in both commodity futures and financial assets. However, properties such as negative asymmetry, positive correlation with stocks indexes and higher volatility levels during the trading session, typical of financial assets, and the existence of inflation hedge and positive correlation with bonds, typical of commodity futures, are also detected. Therefore, our results indicate that EUAs returns do not behave like common commodity futures or financial assets, and point to the fact that EUAs are a new asset class.

Keywords: European Union Allowances (EUAs); Stylised fact; Asset class; Commodity; Commodity prices; Derivatives analysis; Statistics; Time series analysis

JEL Classification: G1

31.1 Introduction

In order to achieve the reduction objectives for greenhouse gas emissions, the European Union decided that a big part of that reduction would have to be directly assumed by the companies of the most polluting sectors. Since 2005, the companies of these sectors, included in the 2003/87/EC Directive, receive each year entitlements to emit one tonne of carbon dioxide equivalent gas, which are denominated European Union Allowances (EUAs).^{*} At the end of the control period, the firms covered by the environmental regulations have to give a sufficient number of allowances to cover their verified real emis-

^{*} The sectors included in the 2003/87/CE Directive are: energy (electricity, co-generation), refining petroleum, iron and steel industry, mineral products, cement, lime, glass, ceramics (roofing tiles, bricks, floor tiles, etc.), cardboard, pulp and paper.

sions. If these companies emit more CO₂ than the allowances they own, they would have to go to the European Union Emission Trading Scheme (EU ETS) and buy the difference. It is important to highlight that not only the polluting companies participate in the EU ETS, but also external agents can trade in it. Therefore, knowledge of the statistical properties of EUAs is of interest not only for hedging operations, but also for speculative or risk management purposes.

In spite of the youth of the EU ETS, the academic literature has analysed the CO₂ market from diverse perspectives. Mansanet-Bataller and Pardo (2008) study the market at an institutional level. Daskalakis and Markellos (2008), Miclaus *et al.* (2008), Daskalakis *et al.* (2009) and Mansanet-Bataller and Pardo (2009) study different aspects of the market efficiency. Benz and Hengelbrock (2008), Rotfuss (2009), Uhrig-Homburg and Wagner (2009) and Rittler (2012) analysed the lead-lag relationship between the spot and futures markets, this last one being the market that leads the price discovery process. Finally, Borak *et al.* (2006), Benz and Trück (2009), Daskalakis *et al.* (2009) and Chesney and Taschini (2012) proposed different alternatives to model the price dynamics of CO₂ emission allowances.*

The papers mentioned above analyse short periods of time and the statistical properties of EUA returns appear as a secondary objective. This is precisely the main purpose of our study. Specifically, starting from the stylised facts described by Cont (2001) for financial assets, and from the statistical properties analysed by Gorton and Rouwenhorst (2006) and Gorton *et al.* (2012) for commodities, we study the stylised empirical facts of EUA returns. By doing so, we will determine whether EUA behaves like a financial asset, a commodity, or a new asset. The knowledge of its statistical properties is of interest for both policy makers and portfolio managers, in terms of regulation and portfolio diversification, respectively. The remainder of the paper is organised as follows. Section 31.2 presents the stylised facts observed in assets grouped into four categories. Section 31.3 describes the data we used in the study. The results of the analysis are presented and discussed in Section 31.4. The last section summarises with some concluding remarks.

31.2 Stylised Facts of Asset Returns

Following Cont (2001, p. 233), we define the stylised facts of an asset as a set of statistical facts which emerge from the empirical study of asset returns and which are common to a large set of assets and markets. Some papers have overviewed the stylised facts that are characteristics of the financial assets (see among others, Pagan 1996, Cont 2001, Morone 2008, and Sewell 2011) and those that are common in the commodities (see Gorton and Rouwenhorst 2006 and Gorton *et al.* 2012). Although some stylised facts are related to each other, the statistical characteristics of asset returns can be grouped into four large sets of phenomena that have to do with the distribution of returns, correlation facts, volatility-related properties and commodity-related facts.

In the first group, the characteristics observed in the historical returns distribution are analysed. The first step is to look into the existence of normality in the distribution of returns. In the case of not rejecting normality, we could not reject both symmetry and the absence of heavy tails. However, the rejection of the hypothesis of normality would make necessary a more exhaustive analysis to determine if the reason for the rejection comes from the existence of asymmetry or because the frequency curve is more or less peaked than the mesokurtic curve. The second aspect we analyse is the *intermittency*, which refers to the phenomenon that returns present, at any time scale, high variability that is translated in the appearance of *outliers* throughout the asset life.[†] Thirdly, we investigate the *aggregational gaussianity*, a third aspect that is detected in historical returns distributions and which makes reference to the fact that the aggregation of data in bigger time intervals approaches the Gaussian data distribution. While the two last phenomena have been observed both in financial assets and commodity futures, the empirical evidence regarding asymmetry usually indicates that the skewness is negative in financial assets and positive in commodities. Finally, we have analysed if EUA returns are stationary. When series

* An excellent review of this kind of literature can be found in Convery (2009).

[†] Outliers refer to those returns which, by their magnitude, are considered as unusual and infrequent.

are non-stationary and follow a unit root process, persistence of shocks will be infinite. If EUA prices follow a trend stationary process, then there exists a tendency for the price level to return to its trend path over time and investors may be able to forecast futures returns by using information on past EUA returns. However, the majority of economic and financial time series exhibit trending behaviour or non-stationarity in the mean. If EUA price series were non-stationary, any shock to EUA price would be permanent, implying that EUA futures returns would be unpredictable based on historical observations.

The second group refers to some correlation facts observed in the returns of whatever asset and to their consequences. Firstly, we study the autocorrelation of the returns. This has been the classical way to test the weak form of the efficient hypothesis in financial markets. The absence of significant linear correlations in returns has been widely documented and it is usually not detected except for very small intraday time scales. Secondly, given that the correlation test may be influenced by the presence of extreme returns, we carry out a test run in order to check the existence of randomness in EUA returns generation. With this test, we get robustness on a possible predictability of EUA returns over the short term. Thirdly, we analyse the presence of a slow decay of autocorrelation in absolute returns, also known as the *Taylor effect*. Following Taylor (2007), this effect makes reference to the fact that a big return in absolute terms is more probably followed by another big one, rather than a small one. This phenomenon, unlike the two previous ones, is interpreted as a sign of long-range dependence.

The third group of facts investigates the specific characteristics observed in volatility. The first well-known property about volatility is the positive autocorrelation observed in different measures of volatility over several days. This fact is generally detected through the existence of autocorrelation in squared returns and it is known as *volatility clustering*. Secondly, we study three volatility-related cross-correlation facts. The first one looks at whether volatility responds differently to positive and negative shocks of the same magnitude, the second one analyses the correlation between volume and volatility, and the last one examines the correlation between the change in the open interest and the volatility.* Although studies on the relationship between trading measures and the underlying price volatility provide mixed evidence, following Bhargava and Malhotra (2007) a number of other researchers report a positive correlation between trading activity and volatility. The significance of these two relationships would indicate that the volume and/or the change in the open interest, in EUA markets, could be used as explanatory variables of volatility.

The last group alludes to specific features extracted from commodity behaviours. The first aspect we investigate is the presence of a non-trading effect observed in the volatility of weather-sensitive assets. Following Fleming *et al.* (2006), trading versus non-trading period variance ratios in weather-sensitive markets are lower than those in the equity markets. Given that Mansanet-Bataller *et al.* (2007) and Alberola *et al.* (2008) have shown empirical evidence about the influence of weather on carbon returns, we study the non-trading effect on EUAs by testing whether the information flow on EUA markets is evenly distributed around the clock. Another stylised fact characteristic of commodities is the negative correlation with stocks and bonds. Generally, commodity futures exhibit a certain negative correlation, mainly in the early part of falling periods, as noted in Gorton and Rouwenhorst (2006). Furthermore, the negative correlation becomes greater as we increase the time lag in which we hold the positions. Finally, we test the property of *inflation hedge*. Assets hedge against inflation when they correlate positively and significantly against it. Following Gorton and Rouwenhorst (2006), commodity futures usually show better behaviour against unexpected inflation than stocks or bonds do, and therefore they can be used for this kind of hedging.

31.3 Data

To carry out the analysis of the stylised facts of EUA returns, we have to select the most representative EUA asset in the market and the time frame. EUAs can be traded in several organised markets such as spot, futures and options markets. From among these markets, most of the trading volume is

* Open interest is the total sum of all outstanding long and short positions of futures contracts that have not been closed.

concentrated in the futures markets, especially in the futures contracts listed for electronic trading at the European Climate Exchange (ECX).^{*} Furthermore, empirical evidence cited in Section 31.1 supports the notion that the price discovery process is led by the futures markets. For these reasons, we have chosen futures prices in order to obtain the most informative EUA return.

Related to the time frame used in this study, it is convenient to clarify some aspects of the EU ETS. Trading in EUAs has been organised into different phases. Phase I covers the years from 2005 to 2007 and was considered as a *pilot or learning* phase, characterised by an excess of EUAs that provoked a sharp decline in prices.[†] Phase II matches the Kyoto protocol fulfilment period, which goes from January 2008 to December 2012, while Phase III will include the period 2013 to 2020.[‡] Another important aspect to consider is that Phase I EUAs could not be used in Phase II, meaning that *banking* was not allowed between Phase I and Phase II. Nevertheless, banking is allowed between Phases II and III, a reason that Phase II EUAs might be used during the 2013 to 2020 period. For these reasons, we limit our sample to Phase II EUA futures returns. In particular, we focus on 2008, 2009 and 2010 EUA ECX futures contracts. The period sample goes from 22 April 2005 to 15 December 2008 for the first contract, from 22 April 2005 to 14 December 2009 for the second one, and from 22 April 2005 to 20 December 2010 for the third one, collecting 936, 1190 and 1446 daily returns for 2008, 2009 and 2010 EUA ECX futures contracts, respectively. Figure 31.1 shows the price evolution for the three EUA




FIGURE 31.1 Prices evolution: the prices evolution is shown for the 2008, 2009 and 2010 EUA futures contracts. The sample period goes from 22 April 2005 to 15 December 2008 for the first contract, from 22 April 2005 to 14 December 2009 for the second one, and from 22 April 2005 to 20 December 2010 for the last one.

^{*} The unit of trading of one contract is one lot of 1000 CO₂ EU allowances. See the User Guide of ICE ECX Contracts: EUAs and CERs at <https://www.theice.com/productguide/ProductGroupHierarchy.shtml?groupDetail=&group.groupId=19> for further information about the contract specifications of ECX EUA futures contract (last accessed on 20 March 2012).

[†] The special features observed in this period are treated in detail by, among others, Miclaus *et al.* (2008), Paoella and Taschini (2008), Benz and Trück (2009), Daskalakis *et al.* (2009) and Mansanet-Bataller and Pardo (2009).

[‡] Given that the EU ETS is organised in phases, the member states must elaborate a National Allocation Plan (NAP) for the first two phases, in which the member states attribute their emission allowances to the different companies, included in the sectors involved in the 2003/87/EC Directive, and establish the emission limits for the different sectors, as well as for each one of the facilities covered by EU ETS. Therefore, these NAPs establish the EUA supply available in the market until 2012. Contrary to Phase I and Phase II, in Phase III only one EU-wide emissions allocation under ETS Phase III will be elaborated, where each installation's individual allocation will be decided by the European Commission.

futures contracts. In general, all three contracts share a common behaviour. Until 24 April 2006, futures prices had a positive trend, over-passing the 30 Euros price per contract, but continuous rumours of over-allocation crashed Phase I prices, negatively affecting Phase II prices too. From then, prices declined until reaching their first minimum of 14 Euros in February 2007. In July 2008, EUA prices approached their maximum quote of over 30 Euros, following which prices started a new bearish period because of the financial crisis. By February 2009, EUA prices reached their historical minimum of around eight Euros. Finally, from April 2009 to the end of 2010, EUA prices have been fluctuating around 14 Euros.

Additionally, in order to test some correlation facts, other series of data have been used: a stock index, a stock index futures contract, bond futures contracts, a commodity futures contract, a risk-free asset series and the observed inflation. We have chosen for all of these series the European benchmarks, using the same time frame as the one we have selected for each EUA contract. Specifically, we have chosen the Euro Stoxx 50 and its futures contract, traded at the EUREX market, as indicative of a European stock index and its future. Regarding interest rates, the fixed income futures have also been obtained from the EUREX market. The three references chosen for this case are Euro Schatz futures, Euro Boll futures, and Euro Bund futures, as benchmarks for European short-, mid- and long-term bond prices, respectively. As a representative commodity futures contract, we have chosen the Brent futures contract traded at the International Petroleum Exchange (IPE). The risk-free asset has been approached through the one-month Euribor rate and, as a proxy of the inflation rate, we have used the European Harmonized Index of Consumer Prices (HICP).

We have carried out our study using returns defined as $r_t = \log(p_t/p_{t-1})$, where p_t is the closing price of the futures contract on day t . Finally, when it has been needed, we have switched futures contracts on the last trading day in order to create a continuous series in futures contracts.

31.4 Results

In this section, we test the presence of the stylised facts described in [Section 31.2](#): the distribution of returns, correlation facts, volatility-related features and commodity-related facts. For each case, we present the methodology used to test the empirical property, the results and the financial implications.

31.4.1 Data Distribution

From a financial point of view, the *gaussianity* of returns data is of interest for portfolio theory, derivatives pricing and risk management. However, one of the most wellknown properties in the distributions of asset returns is precisely the non-normality of the returns series and the fact that they present a greater number of extreme values than those observed in series with Gaussian distributions.* As a background, [Table 31.1](#) presents the summary statistics of the EUA returns for 2008, 2009 and 2010 futures contracts, and [Figure 31.2\(a\)](#), [\(b\)](#) and [\(c\)](#) show the histogram of the historical daily returns distribution. All three daily series are leptokurtic and present negative skewness. The Jarque–Bera statistic rejects the null hypothesis of normality for all three series at the 5% level. The histogram shows heavy tails and a high number of extreme returns, both positive and negative.

The Jarque–Bera statistic is based on the assumption of normality; therefore the daily returns series can be symmetric but non-normal due to the excess of kurtosis. To check this fact, we test the distribution symmetry by applying the method proposed by Peiró (2004). Firstly, we divide each sample into two subsamples, where the first subsample contains the excesses of positive returns with respect to the mean $r^+ = \{r_t - \bar{r} \mid r_t > \bar{r}\}$, and the second one, the excesses of negative returns with respect to the mean, in absolute terms $r^- = \{\bar{r} - r_t \mid r_t < \bar{r}\}$. Secondly, we use the Wilcoxon rank test to check whether the two subsamples come from the same distribution.

* In Catalán and Trávez (2007), some reasons concerning the importance of extreme values are presented.

TABLE 31.1 Summary Statistics for EUA Futures Returns

	Daily	Monday to Monday	Wednesday to Wednesday	Monthly Based Day 1	Monthly Based Day 15	Residuals EGARCH
Panel A: 2008 EUA futures contract						
Mean	-0.0001	-0.0007	-0.0007	-0.0011	-0.0025	-0.0140
Median	0.0013	0.0072	0.0033	-0.0026	0.0022	0.0209
Maximum	0.1865	0.2076	0.1441	0.3189	0.2786	3.9433
Minimum	-0.2882	-0.5146	-0.3757	-0.3780	-0.3405	-7.5034
Std. Dev.	0.0286	0.0767	0.0646	0.1532	0.1317	1.0003
Skewness	-1.3166	-1.8022	-1.3332	-0.2364	-0.2879	-0.5897
Kurtosis	17.9459	13.1149	9.5885	2.7371	3.0530	7.6354
Jarque-Bera	8982.2510	912.8140	397.8340	0.5240	0.5990	891.3000
<i>p</i> -value	0.0000	0.0000	0.0000	0.7690	0.7410	0.0000
K-S stat	2.9130	1.2370	1.1710	0.5710	0.5210	2.6360
<i>p</i> -value	0.0000	0.0940	0.1280	0.9000	0.9490	0.0000
Observations	936	190	189	43	43	935
	Daily	Monday to Monday	Wednesday to Wednesday	Monthly Based Day 1	Monthly Based Day 14	Residuals EGARCH
Panel B: 2009 EUA futures contract						
Mean	-0.0001	-0.0009	-0.0006	-0.0035	-0.0031	-0.0085
Median	0.0007	0.0050	0.0032	-0.0013	0.0132	0.0283
Maximum	0.1932	0.2069	0.1782	0.3219	0.4039	3.6552
Minimum	-0.2811	-0.5087	-0.3667	-0.3697	-0.4580	-7.8527
Std. Dev.	0.0286	0.0757	0.0649	0.1498	0.1487	1.0003
Skewness	-0.8873	-1.4620	-1.0739	-0.2569	-0.3153	-0.5654
Kurtosis	14.4420	10.9417	8.0011	2.7422	4.1379	7.1086
Jarque-Bera	6647.5940	722.1760	2.9740	0.7570	3.8780	899.6650
<i>p</i> -value	0.0000	0.0000	0.0000	0.6850	0.1440	0.0000
K-S stat	2.9180	1.2680	1.3740	0.6240	0.6300	2.9260
<i>p</i> -value	0.0000	0.0800	0.0460	0.8310	0.8220	0.0000
Observations	1190	242	241	55	55	1189
	Daily	Monday to Monday	Wednesday to Wednesday	Monthly Based Day 1	Monthly Based Day 20	Residuals EGARCH
Panel C: 2010 EUA futures contract						
Mean	-0.0001	-0.0008	-0.0005	-0.0015	-0.0030	-0.0003
Median	0.0002	0.0019	0.0018	-0.0020	-0.0020	0.0005
Maximum	0.1912	0.2078	0.1830	0.3249	0.2801	0.1979
Minimum	-0.2743	-0.5016	-0.3582	-0.3654	-0.3710	-0.2711
Std. Dev.	0.0268	0.0708	0.0597	0.1392	0.1254	0.0267
Skewness	-0.8500	-1.4094	-1.0625	-0.2087	-0.2978	-0.8110
Kurtosis	14.5575	11.3714	8.6198	3.1254	3.5760	14.4896
Jarque-Bera	8222.0183	959.0620	442.2060	0.5304	1.9451	8106.6428
<i>p</i> -value	0.0000	0.0000	0.0000	0.7670	0.3781	0.0000
K-S stat	3.0339	1.5029	1.7856	0.4845	0.5570	3.0453
<i>p</i> -value	0.0000	0.0218	0.0034	0.9730	0.9157	0.0000
Observations	1446	295	294	67	68	1445

Notes: This table presents the descriptive statistics for 2008, 2009 and 2010 EUA futures contracts traded on the ICE ECX market. Panel A, B and C present the statistics for the 2008, 2009 and 2010 EUA futures contracts, respectively. The 'Monday to Monday' ('Wednesday to Wednesday') column shows the results for weekly returns calculated from Monday (Wednesday) close to Monday (Wednesday) close. 'Monthly based day 1' presents the results for monthly returns, calculated from the first trading day of each month to the last trading day in the same month. 'Monthly based day 15', 'Monthly based day 14' and 'Monthly based day 20' show the results for monthly returns, calculated from one day in mid-month to the same day in the next month, being the number for that specific day. 'Residuals EGARCH' stands for the statistics of the residuals obtained after fitting an AR(1)-EGARCH(1,1) model for each contract. 'K-S stat' stands for the Kolmogorov-Smirnov statistic, which tests the null hypothesis of the normality of the distribution. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.




FIGURE 31.2 Returns histograms: (a), (b) and (c) depict the histograms of the daily returns for the 2008, 2009 and 2010 EUA futures contracts, respectively, in comparison with the histogram of a normal distribution. The period sample goes from 22 April 2005 to 15 December 2008 for the first contract, from 22 April 2005 to 14 December 2009 for the second one and from 22 April 2005 to 20 December 2010 for the last one, collecting 936, 1190 and 1446 daily returns for 2008, 2009 and 2010 EUA ECX futures contracts, respectively.

Given the negative skewness coefficients reported in Table 31.1, the asymmetry tests, non-reported, show a statistic of 2.9854 (p -value 0.0028) for futures with maturity in 2008, 2.0044 (p -value 0.0450) for 2009 futures and 1.8161 (p -value 0.0694) for 2010 futures. Therefore, we observe a significant negative asymmetry for the first two contracts, something common in financial assets such as stocks, while we cannot reject symmetry at the 5% level for 2010 futures. This characteristic has a direct implication in margins requirements asked by clearing houses in EUA derivatives markets. A negative asymmetry implies more risk for long positions and, as a consequence, bigger margins should be asked for those positions.

The second aspect we have analysed is the *intermittency* of the series. To study the existence of this empirical fact, we have used the VaR approach and, following Galai *et al.* (2008), the procedure based on the calculation of the Huber M-estimator to distinguish between outliers and the body of the distribution. Table 31.2 presents the outlier detection based on daily returns for the EUA 2008, 2009 and 2010 futures contracts. Five percent VaR (95% VaR) provides outlier detection through the Value at Risk method, considering as extreme positive (negative) outliers those returns that are over (under) the 5% one-day VaR (the 95% one-day VaR).^{*} The 5% one-day VaR measure (3.89, 3.98 and 3.78% for 2008, 2009 and 2010 futures, respectively) is of interest for estimating the risk of loss in short positions in futures markets, while the 95% one-day VaR (4.53, 4.56 and 4.24% for 2008, 2009 and 2010 futures, respectively) is relevant for measuring the same risk but for long positions. Obviously, the number of positive and negative extreme returns is the same, but the 95% VaR measure in absolute terms is higher than the 5% VaR measure, confirming the negative asymmetry previously observed.

To get robustness in the outlier detection analysis, we have obtained the Huber M-estimator. The fifth, ninth and thirteenth columns in Table 31.2 present the results following this procedure.[†] Based on the M-estimator, over 17% of returns in all three samples are classified as outliers. Furthermore, in all three samples, the number of negative outliers (90, 110 and 126, respectively) far exceeds the number of positive ones (66, 89 and 105, respectively). Note that in both cases the number of outliers has increased over time.

Panel B of Table 31.2 presents the size of events measured by the number of standard deviations (σ) that an outlier return deviates from the mean. Independently of the approach, the majority of the outliers have a size between one to four standard deviations. This could be considered as usual for normal distributions. However, the empirical VaR and the Huber M-estimator approaches also detect diverse extreme large movements, in all three futures contracts, that diverge from the mean and the median, respectively, from 5 σ s to 11 σ s. Additionally, Panel C of Table 31.2 presents the clusters of outliers. The Huber M-estimator shows concentrations of up to six consecutive outliers while the VaR approach detects clusters of up to three outliers. It is important to note that the number of clusters with this procedure is higher when the outliers are considered independently of their sign, indicating that extreme returns are followed by themselves. On the whole, the results of Table 31.2 indicate a high probability of occurrence of large movements in the EUA market and prove the existence of *intermittency* in EUA returns.

Thirdly, we have studied another aspect that is usually detected in historical returns distribution, which is the fact that the aggregation of data in bigger time intervals approaches a Gaussian data distribution. In order to test the *aggregational gaussianity*, we have generated two subsamples both with weekly and monthly returns. The weekly returns have been calculated taking the close price

^{*} Malevergne and Sornette (2004) summarise the virtues of VaR measures in three characteristics: simplicity, relevance in addressing the ubiquitous large risks often inadequately accounted for by the standard volatility, and their prominent role in the recommendations of the international banking authorities.

[†] To obtain outliers by applying the M-Huber estimator, it is necessary to transform the data until convergence is attained. Furthermore, for the generation of the M-Huber estimator, we have to define the term k , which will be used as a limit during the detection of outliers. Following Galai *et al.* (2008), we have chosen $k = 4.496$. With this selection, the iterative process ends with an M-Huber estimator of 0.000879 for 2008 futures, 0.000561 for 2009 futures and 0.000465 for 2010 futures. A detailed description of this procedure may be found in Hoaglin *et al.* (1983, chapter 11).

TABLE 31.2 Outlier Detection

	2008 EUA Futures Contract				2009 EUA Futures Contract				2010 EUA Futures Contract			
	5% VaR	95% VaR	Total VaR	Huber	5% VaR	95% VaR	Total VaR	Huber	5% VaR	95% VaR	Total VaR	Huber
Number of observations	936				1190				1446			
Panel A: Outliers summary												
Number of outliers	47	46	93	156	60	59	119	199	73	72	145	231
Number of outliers	5.02%	4.91%	9.94%	16.67%	5.04%	4.96%	10.00%	16.72%	5.05%	4.98%	10.03%	15.98%
VaR return	3.89%	-4.53%	—	—	3.98%	-4.56%	—	—	3.78%	-4.24%	—	—
Number of positive outliers	47	0	47	66	60	0	60	89	73	0	73	105
Number of negative outliers	0	46	46	90	0	59	59	110	0	72	72	126
Panel B: Size of outliers (σ s)												
From 0 to 1	0	0	0	0	0	0	0	0	0	0	0	0
From 1 to 2	34	20	54	118	42	28	70	150	48	35	83	170
From 2 to 3	8	18	26	24	11	22	33	33	17	25	42	41
From 3 to 4	4	4	8	9	6	5	11	11	5	8	13	13
From 4 to 5	0	1	1	1	0	1	1	1	2	0	2	2
From 5 to 6	0	2	2	2	0	2	2	2	0	3	3	3
From 6 to 7	1	0	1	1	1	0	1	1	0	0	0	0
From 7 to 8	0	0	0	0	0	0	0	0	1	0	1	1
From 8 to 9	0	0	0	0	0	0	0	0	0	0	0	0
From 9 to 10	0	0	0	0	0	1	1	1	0	0	0	0
From 10 to 11	0	1	1	1	0	0	0	0	0	1	1	1
Panel C: Clusters of outliers												
1 value	37	36	63	84	44	45	73	106	57	54	94	129
2 values	5	5	12	24	8	7	17	31	8	9	18	30
3 values	0	0	2	5	0	0	4	4	0	0	5	6
4 values	0	0	0	1	0	0	0	2	0	0	0	2
5 values	0	0	0	1	0	0	0	1	0	0	0	2
6 values	0	0	0	0	0	0	0	1	0	0	0	1

Notes: This table presents the outlier detection based on daily returns for the EUA 2008, 2009 and 2010 futures contracts. '5% VaR' ('95% VaR') provides the outlier detection through the Value at Risk (VaR) method, considering as extreme positive (negative) outliers those returns which are over (under) the 5% one-day VaR (the 95% one-day VaR). Total VaR considers the aggregation of extreme returns. 'Huber' provides the outlier detection through the M-Huber estimator. The table is divided into three panels. Panel A summarises the outlier detection. Panel B presents the magnitude of the outliers detected in terms of standard deviations of returns (σ). Panel C shows the number of clusters of outliers that have been detected. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

TABLE 31.3 Unit Root Tests

	Level	Differences
2008 EUA futures contract	0.3281	0.0949*
2009 EUA futures contract	0.3867	0.0783*
2010 EUA futures contract	0.3392	0.0733*

Notes: This table summarises the Kwiatkowski *et al.* (1992) unit root tests for the time series of EUAs future prices both in levels and in first differences. We test the null of stationarity with intercept and deterministic time trend for the price series in levels. The critical value (at the 1% level) is 0.216 (see Table 1 in Kwiatkowski *et al.*'s paper, p. 166). If rejected, we then test the null of stationarity plus intercept for the time series in first differences. The critical value is 0.739 (see Table 1 in Kwiatkowski *et al.*'s paper, p. 166). '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

*The null of stationarity cannot be rejected (at the 1% level).

from Monday close to the following Monday close, while the monthly returns take the closing prices of the last trading day of two consecutive months. Table 31.1 presents the statistics for all three subsamples. Given the absence of normality, we have applied the non-parametric Kolmogorov–Smirnov test. Moving from daily returns to lower frequency returns, the normality of the distributions cannot be rejected at the conventional levels of significance. To confirm our results, we have generated two additional subsamples. In the case of weekly returns from Wednesday closing to the following Wednesday closing and for monthly returns calculating returns from two consecutive mid-months days, for all three futures contracts. We obtain similar results, presented in Table 31.1, that confirm the fact that the empirical distribution of EUA returns tends to normality as the frequency of observation decreases.

Finally, we have analysed if EUA returns are stationary by applying the Kwiatkowski *et al.* (1992) unit root test. We have tested the null of stationarity with intercept and deterministic time trend for the price series in levels. The critical value (at the 1% level) is 0.216 (see Table 1 in Kwiatkowski *et al.*'s paper, p. 166). If rejected, we have then tested the null of stationarity plus intercept for the time series in first differences. The results indicate that EUA prices should first be differenced to render the data stationary in all three contracts analysed (see Table 31.3). Therefore EUA prices are difference-stationary and EUA futures returns would be unpredictable based on past observations. This fact should be taken into account for cointegration analysis.

31.4.2 Correlation Facts

The second group of empirical aspects is related to some correlation facts observed in the returns and to their consequences. On the one hand, we study the significance of the autocorrelation coefficients. The study of the presence/absence of autocorrelation is a typical way to detect *short-term predictability of returns* from past information. Given the non-normality of daily distributions, we have decided to apply the non-parametric Spearman's rank correlation coefficient. Table 31.4 presents the Spearman's autocorrelation coefficients and their associated *p*-values. Autocorrelation is significant and positive only for the first lag at the 1% level.

Note that the square of the autocorrelation coefficient could be interpreted as the fraction of the variation of return on day *t* explained by a lagged return in simple linear regression. In the case of one lagged, the variation of today's return is $(0.124)^2 = 1.53$, $(0.109)^2 = 1.10$ and $(0.0779)^2 = 0.61\%$ for 2008, 2009 and

TABLE 31.4 Autocorrelation Tests

Lag	2008 EUA Futures Contract		2009 EUA Futures Contract		2010 EUA Futures Contract	
	ρ	p -Value	ρ	p -Value	ρ	p -Value
1	0.1240	0.0001	0.1090	0.0002	0.0779	0.0031
2	-0.0095	0.7714	-0.0160	0.5814	-0.0250	0.3421
3	0.0444	0.1748	0.0480	0.0979	0.0395	0.1339
4	0.0386	0.2381	0.0484	0.0949	0.0385	0.1434
5	0.0083	0.7996	0.0066	0.8198	0.0131	0.6183
6	-0.0003	0.9936	0.0044	0.8798	0.0016	0.9526
7	-0.0280	0.3922	-0.0170	0.5577	-0.0143	0.5875
8	0.0221	0.4987	0.0156	0.5911	0.0144	0.5858
9	-0.0083	0.7998	0.0073	0.8023	0.0259	0.3257
10	0.0660	0.0435	0.0548	0.0587	0.0398	0.1315

Notes: This table provides the first 10 lagged sample autocorrelation coefficients for the 2008, 2009 and 2010 EUA futures contracts returns. ρ stands for the Spearman's rank correlation coefficient. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

TABLE 31.5 Run Tests

	Observed Number of Runs	Expected Number of Runs	Statistic	p -Value
2008 EUA futures contract	429	469	2.4855	0.0129
2009 EUA futures contract	555	596	1.6820	0.0926
2010 EUA futures contract	698	724	1.3679	0.1797

Notes: This table reports the test run statistics for all three EUA futures contracts where the null hypothesis is that returns are generated in a random way. The table shows the observed and the expected numbers of runs, the statistic and its p -value. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

2010 futures contracts, respectively. Therefore, although the influence of the past return is positive and significant, the short-term predictability is very weak. Additionally, and given that there exist a large number of outliers, we have eliminated their effects in the autocorrelation analysis by applying a test run. This test assumes an ordered sequence of n_1 returns above the median and n_2 returns below the median. A run of type 1 (type 2) is defined as a sequence of one or more returns above (below) the median that are followed and preceded by returns below (above) the median. r_1 and r_2 are the number of runs of type 1 and 2, respectively, and r is the total number of observed runs. By comparing the observed number of runs and the expected number of runs, we test the hypothesis that returns follow a random sequence. Too few or too many runs would suggest non-randomness in the distribution. Table 31.5 shows results that indicate that the null of randomness is rejected at the 5% level for 2008 futures, but not for 2009 and 2010 futures contracts. In summary, correlation analysis and test runs indicate a positive but weak short-run predictability of past returns. Furthermore, this small predictability has diminished over time.

Finally, we analyse the presence of a slow decay of autocorrelation in absolute returns, also known as the Taylor effect, which would indicate that large returns in absolute terms are more probably followed by another large return, rather than a small one. Note that we have already obtained preliminary evidence of dependence among outliers; however, the Taylor effect makes reference to a feature that is present in the entire sample. To test this effect, we have obtained the Spearman's rank autocorrelation coefficient taking into account absolute returns. Table 31.6 presents results that indicate a significant slow decay of sample autocorrelation coefficients at the 5% level, suggesting a long-range dependence in absolute returns.

TABLE 31.6 Taylor Effect

Lag	2008 EUA Futures Contract		2009 EUA Futures Contract		2010 EUA Futures Contract	
	ρ	p -Value	ρ	p -Value	ρ	p -Value
1	0.1071	0.0010	0.1322	0.0000	0.1301	0.0000
2	0.0498	0.1280	0.0848	0.0034	0.1106	0.0000
3	0.1202	0.0002	0.1533	0.0000	0.1583	0.0000
4	0.0792	0.0154	0.1012	0.0005	0.0964	0.0002
5	0.0873	0.0075	0.1156	0.0001	0.1233	0.0000
6	0.0907	0.0055	0.1331	0.0000	0.1398	0.0000
7	0.0639	0.0507	0.0722	0.0127	0.0776	0.0032
8	0.0740	0.0235	0.1153	0.0001	0.1268	0.0000
9	0.0861	0.0084	0.1102	0.0001	0.0988	0.0002
10	0.0466	0.1541	0.0758	0.0089	0.0874	0.0009

Notes: This table provides the first 10 lagged sample autocorrelation coefficients for the 2008, 2009 and 2010 EUA futures contracts absolute returns. ρ stands for the Spearman's rank correlation coefficient. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

31.4.3 Volatility-Related Features

The third group of facts investigates some aspects of EUA returns related to some phenomena of the time series volatility. The first is related to the positive autocorrelation observed in different measures of volatility; the second makes reference to the fact that volatility responds differently to positive and negative shocks of the same magnitude; and the last one analyses the correlation between trading-related measures and volatility.

Firstly, we have depicted the squared returns of all three futures contracts as a standard approach to detect volatility clusters. Figure 31.3(a), (c) and (e) exhibit this fact, that is very common both in financial and commodities assets series.

Additionally, we have obtained a high–low daily volatility measure. We have chosen the volatility measure proposed by Garman and Klass (1980) due to its higher relative efficiency than the standard estimators. The Garman and Klass volatility measure (hereinafter, the Garman–Klass volatility) is obtained in the following way:

$$\sigma_{GK} = 0.511(u - d)^2 - 0.019\{c(u + d) - 2ud\} - 0.383c^2,$$

where $u = \log H_t - \log O_t$ is the normalised high (H_t is today's high price; O_t is today's opening price); $d = \log L_t - \log O_t$ is the normalised low (L_t is today's low price); and $c = \log C_t - \log O_t$ is the normalised close (C_t is today's close price).

Figure 31.3(b), and (f) represent the evolution of the Garman–Klass volatility over the trading days of all three futures. Both figures show *volatility clustering*. There are two striking clusters in July 2005 and April 2006. Benz and Hengelbrock (2008) relate the first one to unexpected selling by some Eastern European countries and to the insecurity perceived in the market as a consequence of the terrorist attacks in London in July 2005. The second cluster took place in April 2006 and is explained by the fact that in April 2006 the market understood that there was a great excess of allowances in Phase I, which also affected Phase II EUA futures.

Panel A of Table 31.7 presents the positive and significant Spearman's cross autocorrelation coefficients of Garman–Klass volatility for the first 10 lags. The results confirm the existence of volatility clusters that decay very slowly.




FIGURE 31.3 Volatility clustering: (a), (c) and (e) show the squared returns representation for the 2008, 2009 and 2010 EUA futures contracts, respectively; (b), (d) and (f) show the Garman-Klass volatility representation for the 2008, 2009 and 2010 EUA futures contracts, respectively. Volatility clustering is observed in all figures. ‘2008 EUA futures contract’ refers to the futures contract maturing on 15 December 2008, ‘2009 EUA futures contract’ refers to the futures contract maturing on 14 December 2009 and ‘2010 EUA futures contract’ refers to the futures contract maturing on 20 December 2010.




FIGURE 31.3 (Continued) Volatility clustering: (a), (c) and (e) show the squared returns representation for the 2008, 2009 and 2010 EUA futures contracts, respectively; (b), (d) and (f) show the Garman–Klass volatility representation for the 2008, 2009 and 2010 EUA futures contracts, respectively. Volatility clustering is observed in all figures. ‘2008 EUA futures contract’ refers to the futures contract maturing on 15 December 2008, ‘2009 EUA futures contract’ refers to the futures contract maturing on 14 December 2009 and ‘2010 EUA futures contract’ refers to the futures contract maturing on 20 December 2010.

TABLE 31.7 Volatility-Related Features

Lag	2008 EUA Futures Contract		2009 EUA Futures Contract		2010 EUA Futures Contract	
	Correlation	<i>p</i> -Value	Correlation	<i>p</i> -Value	Correlation	<i>p</i> -Value
Panel A: $\rho_s(\sigma, \sigma_{-\tau})$						
1	0.7634	0.0000	0.8428	0.0000	0.8641	0.0000
2	0.7132	0.0000	0.8384	0.0000	0.8514	0.0000
3	0.6968	0.0000	0.8267	0.0000	0.8400	0.0000
4	0.7207	0.0000	0.8304	0.0000	0.8287	0.0000
5	0.7189	0.0000	0.8266	0.0000	0.8377	0.0000
6	0.7036	0.0000	0.8090	0.0000	0.8249	0.0000
7	0.6796	0.0000	0.8128	0.0000	0.8243	0.0000
8	0.6515	0.0000	0.8034	0.0000	0.8218	0.0000
9	0.6651	0.0000	0.8025	0.0000	0.8181	0.0000
10	0.6592	0.0000	0.8156	0.0000	0.8226	0.0000
Panel B: $\rho_s(\sigma, r_{-\tau})$						
0	-0.1008	0.0020	-0.0931	0.0013	-0.0666	0.0113
1	-0.0915	0.0051	-0.0781	0.0070	-0.0564	0.0321
2	-0.0777	0.0174	-0.0608	0.0361	-0.0485	0.0652
3	-0.0626	0.0557	-0.0777	0.0074	-0.0637	0.0154
4	-0.0655	0.0453	-0.0614	0.0341	-0.0572	0.0298
5	-0.0613	0.0609	-0.0387	0.1821	-0.0580	0.0278
6	-0.0692	0.0342	-0.0630	0.0298	-0.0571	0.0304
7	-0.0506	0.1222	-0.0510	0.0787	-0.0341	0.1956
8	-0.0554	0.0901	-0.0497	0.0867	-0.0541	0.0402
9	-0.0572	0.0803	-0.0312	0.2815	-0.0497	0.0594
10	0.0789	0.0158	-0.0676	0.0197	-0.0589	0.0256

Notes: Panel A presents the Spearman's autocorrelation coefficients (ρ) between the daily Garman–Klass volatility (σ) and the daily Garman–Klass volatility lagged τ periods. Panel B presents the Spearman's cross correlations coefficients between the daily Garman–Klass volatility and the daily futures return (r) lagged τ periods. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

Secondly, we study the *asymmetric volatility*, which is the fact that volatility responds differently to positive and negative shocks of the same magnitude.* Cont (2001), Bouchaud *et al.* (2007) and Bouchaud and Potters (2001) comment on the existence of a negative correlation between volatility and returns in financial assets, and particularly in stocks, which tends to zero as we increase the time lag. Panel B in Table 31.7 present the Spearman's cross correlation coefficients between Garman–Klass volatility and returns. Both the contemporaneous and the one-period lagged correlation coefficient are significant and negative at the 1% level, confirming that EUA volatility responds differently to positive and negative returns.

Although volatility modelling is not an objective of this study, to get additional insights about volatility asymmetry we have tested standard GARCH models that take into account this stylised fact in Phase II.† Following Rodríguez and Ruiz (2009), we have estimated EGARCH models as the most flexible GARCH models that capture the asymmetric effect. In all three series of data, the asymmetric parameter is significant at the 1% significance level; however, the residuals of the estimated models still show persistence of heavy tails (see the last column in Table 31.1). Therefore, when modelling EUA volatility, we suggest assuming fat-tailed unconditional distributions in order to capture all the volatility features of EUA return data.‡

* A survey of the empirical literature about this fact can be found in Taylor (2005).

† Borak *et al.* (2006), Miclaux *et al.* (2008), Paoletta and Taschini (2008) and Benz and Trück (2009), among others, present GARCH models to analyse the conditional volatility of Phase I EUAs returns.

‡ See Rachev and Mitnik (2000, chapter 4) for a comparison of conditional and unconditional distributional models for returns on the Nikkei 225 stock market index.

TABLE 31.8 Trading-Related Features

Lag	2008 EUA Futures Contract		2009 EUA Futures Contract		2010 EUA Futures Contract	
	Correlation	<i>p</i> -Value	Correlation	<i>p</i> -Value	Correlation	<i>p</i> -Value
Panel A: $\rho_s(\sigma, \text{VOL}_{-\tau})$						
0	0.7644	0.0000	0.8832	0.0000	0.8403	0.0000
1	0.7070	0.0000	0.8559	0.0000	0.8210	0.0000
2	0.6875	0.0000	0.8492	0.0000	0.8122	0.0000
3	0.6865	0.0000	0.8512	0.0000	0.8132	0.0000
4	0.6980	0.0000	0.8454	0.0000	0.8082	0.0000
5	0.6956	0.0000	0.8444	0.0000	0.8082	0.0000
6	0.6831	0.0000	0.8377	0.0000	0.8049	0.0000
7	0.6758	0.0000	0.8354	0.0000	0.8045	0.0000
8	0.6667	0.0000	0.8298	0.0000	0.7981	0.0000
9	0.6710	0.0000	0.8352	0.0000	0.7918	0.0000
10	0.6690	0.0000	0.8390	0.0000	0.7946	0.0000
Panel B: $\rho_s(\sigma, \Delta \text{OI}_{-\tau})$						
0	0.1720	0.0000	0.3186	0.0000	0.2705	0.0000
1	0.1146	0.0006	0.2767	0.0000	0.2658	0.0000
2	0.1025	0.0021	0.2817	0.0000	0.2460	0.0000
3	0.1319	0.0001	0.2934	0.0000	0.2536	0.0000
4	0.1286	0.0001	0.2977	0.0000	0.2554	0.0000
5	0.1087	0.0011	0.2793	0.0000	0.2570	0.0000
6	0.1019	0.0022	0.2581	0.0000	0.2641	0.0000
7	0.1148	0.0006	0.2581	0.0000	0.2504	0.0000
8	0.0841	0.0117	0.2588	0.0000	0.2419	0.0000
9	0.1203	0.0003	0.2838	0.0000	0.2325	0.0000
10	0.1394	0.0000	0.2649	0.0000	0.2334	0.0000

Notes: Panel A presents the Spearman's cross correlation coefficients (ρ) between the daily Garman–Klass volatility (σ) and the daily trading volume (VOL) lagged τ periods. Panel B presents the Spearman's cross correlation coefficients between the daily Garman–Klass volatility (σ) and the change in the daily open interest (ΔOI) lagged τ periods. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

Finally, we have tested the relationship between trading activity and volatility. Given that our measure of volatility is an intraday volatility, we have chosen two flow variables, the level of daily trading volume (VOL) and the change in the open interest (ΔOI), to calculate the cross correlations between these trading-related measures and volatility.

Panels A and B in Table 31.8 present results for the cases of volume and change in open interest, respectively. The contemporaneous cross correlation coefficients in the case of volume are positive and significant (76, 88 and 84% for 2008, 2009 and 2010 futures contract, respectively). The same occurs in the case of open interest (17, 32 and 27% for 2008, 2009 and 2010 futures contracts, respectively). Furthermore, the cross correlations for lags from one to ten decay very slowly in both measures and in all three futures contracts. All in all, these facts indicate that both trading-related measures should be considered as explanatory variables in volatility models. Additionally, following the classical approach in the financial literature, which considers volume as a proxy for day trading and speculative activity, and open interest as a proxy to measure hedging activity, we could say that both speculators and hedgers destabilise the EUA market.

TABLE 31.9 Information Flows

	No. observations	σ_{OC}	σ_{CO}	BF-Statistic	<i>p</i> -Value
Panel A: 2008 EUA futures contract					
Weekend	144	47.43	25.17	11.3814	0.0008
Weekdays	574	41.01	28.24	30.0853	0.0000
BF-statistic		0.0144	0.7702		
<i>p</i> -value		0.9044	0.3805		
Total	718	42.36	27.66	41.2757	0.0000
Panel B: 2009 EUA futures contract					
Weekend	142	42.14	34.53	2.0256	0.1558
Weekdays	572	39.88	28.14	25.2061	0.0000
BF-statistic		0.6069	6.3929		
<i>p</i> -value		0.4362	0.0117		
Total	714	40.32	29.57	25.5893	0.0000
Panel C: 2010 EUA futures contract					
Weekend	161	34.21	25.17	10.7123	0.0012
Weekdays	656	30.76	24.91	32.4802	0.0000
BF-statistic		1.1193	0.0713		
<i>p</i> -value		0.2904	0.7895		
Total	817	31.49	24.95	43.1415	0.0000

Notes: This table reports the Brown–Forsythe statistic (BF-statistic), which tests the null of equality of variances among trading and non-trading periods. Panels A, B and C present the results for weekends and for weekdays for 2008, 2009 and 2010 futures contracts, respectively. Weekend results are based on returns from Friday close to Monday close and weekday results are based on returns from Monday close to Friday close. The table also reports the number of observations and the annualised volatility as a percentage for trading (σ_{OC}) and non-trading (σ_{CO}) periods. ‘2008 EUA futures contract’ refers to the futures contract maturing on 15 December 2008, ‘2009 EUA futures contract’ refers to the futures contract maturing on 14 December 2009 and ‘2010 EUA futures contract’ refers to the futures contract maturing on 20 December 2010.

31.4.4 Commodity features

The last group of relevant stylised facts is focused on specific features observed in commodities series. The first we study is the non-trading effect in EUA volatility in order to test whether the information flow is more evenly distributed around the clock in weather-sensitive markets than in the equity market. Firstly, we have tested a possible daily seasonality in returns or variances. Neither non-parametric Kruskal–Wallis nor Brown–Forsythe tests can reject the equality of medians and variances.* Therefore, returns and variances are evenly distributed over the trading days. Secondly, two subsamples have been generated for each maturity. We have separated the weekend from the rest of the week, with the returns from Friday closing to Monday closing standing for weekend returns. Finally, we have split each subsample into two, separating between trading periods (open-to-close returns, OC) and non-trading periods (close-to-open returns, CO).

The results are presented in Table 31.9. All three futures contracts show similar results. We can reject the equality in variances between weekend and weekdays neither for trading nor non-trading periods, with the exception of non-trading periods for 2009 futures contract. However, when comparing volatility between trading and non-trading periods conditioned on weekend or weekdays, we reject the null hypothesis at the 1% level of significance (except for weekend periods for 2009 futures contract), showing that the trading volatility (OC) is higher than non-trading volatility (CO). Taking into account that the number of trading hours (7:00 to 17:00, London local time) is lower than the number of non-trading

* The Kruskal–Wallis statistic is 3.849 (*p*-value = 0.427) for 2008 futures, 4.210 (*p*-value = 0.378) for 2009 futures and 6.4956 (*p*-value = 0.1651) for 2010 futures. The Brown–Forsythe statistic takes the values 1.591 (*p*-value = 0.174), 1.817 (*p*-value = 0.123) and 1.9257 (*p*-value = 0.1037) for 2008, 2009 and 2010 futures contracts, respectively.

TABLE 31.10 Correlation with bonds, stocks and Brent

Asset	2008 EUA Futures Contract		2009 EUA Futures Contract		2010 EUA Futures Contract	
	ρ	p -Value	ρ	p -Value	ρ	p -Value
Schatz bond futures	-0.0950	0.0038	-0.1311	0.0000	-0.1145	0.0000
Bolb bond futures	-0.0971	0.0031	-0.1442	0.0000	-0.1248	0.0000
Bund bond futures	-0.0951	0.0038	-0.1465	0.0000	-0.1268	0.0000
Euro Stoxx 50 futures	0.0585	0.0751	0.1476	0.0000	0.1412	0.0000
Euro Stoxx 50 index	0.0577	0.0787	0.1434	0.0000	0.1362	0.0000
Brent oil futures	0.2413	0.0000	0.2499	0.0000	0.2302	0.0000

Notes: The table gives the Spearman's rank correlation coefficients and their p -values between the EUA 2008, 2009 and 2010 futures contracts and Schatz bond futures (short-term fixed rent), Bolb bond futures (medium-term fixed rent), Bund bond futures (long-term fixed rent), Euro Stoxx 50 futures, the Euro Stoxx 50 index, and the Brent futures contract. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

hours, our results clearly indicate that information flow is concentrated during the business day, similar to the equity markets, and does not evolve randomly around the clock.

A second phenomenon characteristic of commodities futures is the negative correlation with stocks and bonds (see Gorton and Rouwenhorst 2006). We have analysed this aspect by calculating the existing correlation between EUA futures and the main European benchmarks for bond and stock markets. Specifically, we have chosen Schatz bond futures, Bolb bond futures and Bund bond futures as representatives of short-term, medium-term and long-term bonds, and the Euro Stoxx 50 futures contract and its underlying asset as a benchmark for the stock market. Table 31.10 presents Spearman's rank correlation coefficients between the different described assets and the three EUA futures contracts.

The significance of the correlation with all the assets can be seen. Nevertheless, the correlation is negative with bond assets and positive with stocks and Brent, mainly in the 2009 and 2010 futures contracts. A possible explanation of these results can be found in the fact that, in business cycle expansions (contractions), when the stock market is normally in a bullish (bearish) moment, companies produce more (less) and as a consequence they pollute more (less) and need higher (lower) numbers of EUAs. Expansions (contractions) are normally accompanied by an increase (decrease) in interest rates and a decrease (increase) in bond prices. Note that the whole significant correlation picture has appealing implications for portfolio diversification purposes.

Finally, we test the property of *inflation hedge*. If the correlation between EUAs and inflation were significant and positive, they would let us hedge against inflation. To study this fact, we have carried out a monthly analysis. Furthermore, inflation has been divided into three types: the observed, the expected and the unexpected. It is interesting to extract the unexpected component of inflation because this is the part we are interested in hedging against since the expected component is anticipated by the interest rate, and, with fixed income assets we might hedge against it. The Harmonized Index of Consumer Prices (HICP) monthly series has been chosen for the observed inflation.* The expected inflation is the free risk interest rate, which is the reason we use the one month Euribor rate as a proxy, leaving the non-expected inflation as the difference between the observed and the expected inflation.

* Eurostat defines the Harmonized Index of Consumer Prices (HICP) as an economic indicator constructed to measure the changes over time in the prices of consumer goods and services acquired by households. The HICP gives comparable measures of inflation in the euro-zone, the EU, the European Economic Area and for other countries including accession and candidate countries. They are calculated according to a harmonised approach and a single set of definitions. They provide the official measure of consumer price inflation in the euro-zone for the purposes of monetary policy in the euro area and assessing inflation convergence as required under the Maastricht criteria. The elected series is the European Union HICP with 25 countries. More details about this index can be found at <http://ec.europa.eu/eurostat> (last accessed on 20 March 2012).

TABLE 31.11 Inflation hedge

	Correlation		
	Observed Inflation	Expected Inflation	Non-Expected Inflation
<i>2008 EUA futures contract</i>			
Correlation coefficient	0.2895	−0.0908	0.3241
Sig. (bilateral)	0.0601	0.5627	0.0345
<i>2009 EUA futures contract</i>			
Correlation coefficient	0.2953	−0.0618	0.2863
Sig. (bilateral)	0.0290	0.6542	0.0345
<i>2010 EUA futures contract</i>			
Correlation coefficient	0.3379	−0.0316	0.2973
Sig. (bilateral)	0.0054	0.7998	0.0149

Notes: The table shows the Spearman's rank correlation coefficients and their *p*-values between the EUA futures contracts returns and the series of the observed, expected and non-expected inflation. The number of monthly observations is 43, 55 and 67 for the 2008, 2009 and 2010 EUA futures contracts, respectively. '2008 EUA futures contract' refers to the futures contract maturing on 15 December 2008, '2009 EUA futures contract' refers to the futures contract maturing on 14 December 2009 and '2010 EUA futures contract' refers to the futures contract maturing on 20 December 2010.

Table 31.11 shows the correlation coefficients between our asset and the different definitions of inflation. Note that the Spearman's correlation coefficients between the futures and the unexpected inflation are significant in all the futures contracts; therefore, as commodity futures, EUA assets could help us to hedge against inflation.

31.5 Conclusions

The objective of our paper is to study the stylised facts of EUA returns. Given that Phase I (2005–2007) is generally considered as a pilot and learning phase, we have chosen 2008, 2009 and 2010 futures contracts to test the accomplishment of several statistical features that usually appear in both financial and commodities assets.

We have found evidence of intermittency, *aggregational gaussianity*, short-term predictability, the Taylor effect, volatility clustering, asymmetric volatility, and a positive relation between volatility and volume and between volatility and change in open interest. Our findings suggest that temporal dependences should be modelled with ARMA-GARCH structures. However, the persistence of conditional heavy-tails indicates that, in order to capture the volatility features of EUA return data; fattailed unconditional distributions should be assumed. Furthermore, unlike commodities properties, we have observed negative asymmetry, positive correlation with stocks indexes and higher volatility levels during the trading session—an indication that EUA information flow is concentrated during the trading day and does not evolve randomly around the clock. The property of inflation hedge and the positive correlation with bonds, both characteristics being typical of commodity futures, are also detected. Therefore, our results indicate that EUAs do not behave like either common commodities futures or financial assets, and suggest that the EUA is a new asset class. The entirety of these facts, robust over time, has appealing implications for portfolio analysis, volatility modelling, hedging activities and cointegration analysis.

Acknowledgements

A preliminary version of this paper was published by Instituto Valenciano de Investigaciones Económicas (IVIE) as Working Paper WP-AD 2012-14. The authors are grateful for the financial support of the Spanish Ministry of Education and Science (projects ECO2009- 14457-C04-04 and CGL2009-09604), the Fonds Européen de Développement Régional (FEDER), and the Cátedra de Finanzas

Internacionales—Banco Santander. They are also indebted to the European Climate Exchange (ECX) market for providing the database. The usual caveats apply.

References

- Alberola, E., Chevallier, J. and Chèze, B., Price drivers and structural breaks in European carbon prices 2005–2007. *Energy Policy*, 2008, **36**, 787–797.
- Benz, E. and Hengelbrock, J., Liquidity and price discovery in the European CO₂ futures market: An intra-day analysis. Working Paper, Bonn Graduate School of Business, 2008.
- Benz, E. and Trück, S., Modelling the price dynamics of CO₂ emission allowances. *Energy Econ.*, 2009, **31**, 4–15.
- Bhargava, V. and Malhotra, D.K., The relationship between futures trading activity and exchange rate volatility, revisited. *J. Multinat. Finan. Mgmt*, 2007, **17**, 95–111.
- Borak, S., Härdle, W., Trück, S. and Weron R., Convenience yields for CO₂ emission allowances futures contracts. Discussion Paper 2006-076, SFB 649 Economic Risk, Humboldt-University of Berlin, 2006.
- Bouchaud, J.P., Matacz, A. and Potters, M., The leverage effect in financial markets: Retarded volatility and market panic. *Phys. Rev. Lett.*, 2007, **87**, 228701.
- Bouchaud, J.P. and Potters, M., More stylized facts of financial markets: Leverage effect and downside correlations. *Physica A*, 2001, **299**, 60–70.
- Catalán, B. and Trávez, F.J., Forecasting volatility in GARCH models with additive outliers. *Quant. Finan.*, 2007, **7**, 591–596.
- Chesney, M. and Taschini, L., The endogenous price dynamics of emission allowances: An application to CO₂ option pricing. *Appl. Math. Finan.*, 2012, forthcoming.
- Cont, R., Empirical properties of asset returns: Stylized facts and statistical issues. *Quant. Finan.*, 2001, **1**, 223–236.
- Convery, F.J., Reflections—the emerging literature on emission trading in Europe. *Rev. Environ. Econ. & Policy*, 2009, **3**, 121–137.
- Daskalakis, G. and Markellos, R.N., Are the European carbon markets efficient? *Rev. Futures Mkts*, 2008, **17**, 103–128.
- Daskalakis, G., Psychoyios, D. and Markellos, R.N., Modelling CO₂ emission allowances prices and derivatives: Evidence from the European trading scheme. *J. Banking & Finance*, 2009, **33**, 1230–1241.
- Directive 2003/87/EC of the European Parliament and of the Council of 13 October 2003 establishing a scheme for greenhouse gas emission allowance trading within the Community and amending Council Directive 96/61/EC. Available online at: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2003:275:0032:0032:EN:PD> (accessed 1 December 2011).
- Fleming, J., Kirby, C. and Ostdiek, B., Information, trading, and volatility: Evidence from weather-sensitive markets. *J. Finan.*, 2006, **61**, 2899–2930.
- Galai, D., Kedar-Levy, H. and Schreiber, B.Z., Seasonality in outliers of daily stock returns: A tail that wags the dog? *Int. Rev. Finan. Anal.*, 2008, **17**, 784–792.
- Garman, M.B. and Klass, M.J., On the estimation of security price volatilities from historical data. *J. Bus.*, 1980, **53**, 67–78.
- Gorton, G.B., Hayashi, F. and Rouwenhorst K.G., The fundamentals of commodity futures returns. *Rev. Finan.*, 2012, forthcoming.
- Gorton, G.B. and Rouwenhorst, K.G., Facts and fantasies about commodity futures. *Finan. Analysts J.*, 2006, **62**, 47–68.
- Hoaglin, D.C., Mosteller, F. and Tukey, J.W., *Understanding Robust and Exploratory Data Analysis*, 1983 (Wiley Classics Library: New York).
- Kwiatkowski, D., Phillips, P., Schmidt, P. and Shin, Y., Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *J. Econometr.*, 1992, **54**, 1–3, 159–178.

- Malevergne, Y. and Sornette, D., Value-at-Risk-efficient portfolios for a class of super- and subexponentially decaying assets return distributions. *Quant. Finan.*, 2004, **4**, 17–36.
- Mansanet-Bataller, M. and Pardo, Á., What you should know about carbon markets. *Energies*, 2008, **1**, 120–153.
- Mansanet-Bataller, M. and Pardo, Á., Impacts of regulatory announcements on CO₂ prices. *J. Energy Mkts*, 2009, **2**, 77–109.
- Mansanet-Bataller, M., Pardo, Á. and Valor, E., CO₂ prices, energy and weather. *Energy J.*, 2007, **28**, 73–92.
- Miclaus, P.G., Lupu, R., Dumitrescu, S.A. and Bobirca, A., Testing the efficiency of the European carbon futures market using event-study methodology. *Int. J. Energy & Environ.*, 2008, **2**, 121–128.
- Morone, A., Financial markets in the laboratory: An experimental analysis of some stylized facts. *Quant. Finan.*, 2008, **8**, 513–532.
- Pagan, A., The econometrics of financial markets. *J. Emp. Finan.*, 1996, **3**, 15–102.
- Paoletta, M.S. and Taschini, L., An econometric analysis of emission allowances prices. *J. Banking & Finan.*, 2008, **32**, 2022–2032.
- Peiró, A., Asymmetries and tails in stock index returns: Are their distributions really asymmetric? *Quant. Finan.*, 2004, **4**, 37–44.
- Rachev, S. and Mittnik, S., *Stable Paretian Models in Finance*, 2000 (Wiley: Chichester).
- Rittler, D., Price discovery and volatility spillovers in the European Union Emissions Trading Scheme: A high-frequency analysis. *J. Banking & Finan.*, 2012, **36**, 774–785.
- Rodríguez, M.J. and Ruiz, E., GARCH models with leverage effect: Differences and similarities. UC3M Working Paper, Statistics and Econometrics 09-02, Instituto Valenciano de Investigaciones Económicas, 2009.
- Rotfuss, W., Intraday price formation and volatility in the European Union Trading Scheme: An introductory analysis. Discussion Paper 09-018, Centre for European Economic Research (ZEW), Mannheim, Germany, 2009.
- Sewell, M., Characterization of financial time series. Research Note, University College London, 2011.
- Taylor, S.J., *Asset Price Dynamics, Volatility and Prediction*, 2005 (Princeton University Press: Princeton, NJ).
- Taylor, S.J., *Modelling Financial Time Series*, 2nd ed., 2007 (World Scientific: Hackensack, NJ).
- Uhrig-Homburg, M. and Wagner, M., Futures price dynamics of CO₂ emission allowances: An empirical analysis of the trial period. *J. Derivatives*, 2009, **17**, 73–88.

Index

A

Abnormal return of wine index, 179
Aggregational gaussianity, 668, 674
Aggressive monetary policy, 479
Agricultural commodities, 482–485
AIC, *see* Akaike Information Criterion
Air Transportation sector, 364, 379, 382
Akaike Information Criterion (AIC), 62–63, 107, 307, 367, 511
 α -Stable distribution, 552, 557
 estimation results for, 559–561
Aluminium, industrial metals, 476, 479
Analysis of variance, 222
ARCH-LM test, 287–288
Asset prices, price discovery impact measurement on, 471
Asset returns
 data distribution of, 671–676
 descriptive statistics of, 231
 model, 59–60
 stylised facts of, 668–669
Asymmetric cointegration approach, 181
Asymmetric dependence, 413
Asymmetric Exponential Power (AEP), 557–558
 distribution laws, 552
 estimation results for, 559–561
Asymmetric volatility, 681
Asymmetry hypotheses, 368
Asymmetry, multivariate conditional copula with, 407–408
 copula function model, 409–410
 model estimation and portfolio optimization, 410–412
 univariate processes model, 408–409
Asymmetry, tests of, 369–376
Asymptotic Chi-squared distribution, 367
Augmented Dickey–Fuller (ADF) test, 74, 100, 181
Autocorrelation coefficient, 676–677
Autocorrelation of returns, 669
Autoregressive model, 236

AutoRegressive Moving Average (ARMA)

 model, 218, 221

 techniques, 214

Average daily return of wine indexes, 177

B

Backpropagation (BP) neural network
 impact of commission, 294–297
 prediction accuracy comparison, 290–292
 rate of return comparison, 292–293
Backwardation, 4
 pattern, 440
Backward integration method, 581, 582
Balanced stepdown procedure, 120–121, 127–129, 134
Bates model, 634
Bayesian simulation-based estimation techniques, 505
Bayes rule, 505
BEKK-GARCH model, 232
Bernoulli random variables, 505
Bid stack
 choice of parameters, 655
 concrete two-fuel model, 653
 electricity markets price setting, 647–648
Bilateral exchange rates, 231, 235–236
Binomial tree algorithm, 608
Black box price forecasting methods, 10
Black–Scholes–Merton options pricing formula, 502
Black–Scholes (BS) models, 152, 156, 162, 163, 166, 586, 597
 swing options prices in, 618–619
Black’s formula, 517, 528
Bloomberg survey data-set, 472
Bond market, 441
 implied short rate from, 442
 optimal proportions for, 446, 447
Bond-stock mix allocation, 434
Bootstrapping techniques, 125
BP neural network, *see* Backpropagation neural network
Brent blend crude oil (ITCO), 87–90
Brown–Forsythe tests, 683

Brownian motion, 102, 435, 436, 455, 504, 525, 530, 534,
542, 593, 601, 636, 638
hedging cost, 640
BS models, *see* Black–Scholes models
Building Construction sector, 378
Burr III distribution, 153
Business cycle
commodity markets relation to, 488–493
measurement of, 487–488
stock returns and, 305–309
variables, 20

C

CAC40 index, 176, 184
Calendar spread options (CSOs), 7
Calendar spreads, 72, 121
Calibration model
commodity and interest market for, 535
crack spread: heating oil/WTI crude oil, 84–86
location spread: Brent/WTI crude oil, 88–89
long-term process, 576–577
of option pricing models, inverse theory in, 152,
154–155
price forward curve, 573–574
process, 533
short-term process, 574–575
spread process, 79–81
term structure of volatility, 577–578
with time-dependent volatilities, 527–530
Call option, volatility surface, 579, 580
Call spread options, 543, 544
Calmar index, 461
Cap-and-trade market, 656–657
Cap-and-trade regulation structured, 648
Cap-and-trade scheme, 647
choice of parameters, 656
Cap-and-trade systems *vs.* carbon tax, 660–662
Capital Asset Pricing Model, 175
Capital market theory, 338
CAPM model, 318
realized *vs.* expected returns, 323–326
Carbon-di-oxide (CO₂) emissions, 646
Carbon emission, 647
Carbon tax, cap-and-trade systems *vs.*, 660–662
Catalytic cracking method, 99
CAViaR model, *see* Conditional Autoregressive Value
at Risk
CBOT, *see* Chicago Board of Trade
CCAPM model, 318
realized *vs.* expected returns, 323, 327–329
CCC, *see* Constant conditional correlation
Central Limit Theorem, 556
Centre for Research in Security Prices (CRSP)
database, 362
Certainty equivalent, computation of, 454–456
CFNAI, *see* Chicago Fed National Activity Index

CFTC, *see* Commodity Futures Trading Commission
CGMY model, swing options prices in, 618–619
Chemical, Plastic and Transportation Equipment
sectors, 379, 380
Chemical sector, 380
Chicago Board of Trade (CBOT), 230
soybean futures, 191–194, *see also* Soybean futures
Chicago Fed National Activity Index (CFNAI), 304
Chicago Mercantile Exchange (CME), 215, 440
China
metal futures market in, 280
soybeans in, 192–193
Chi-squared likelihood ratio test statistic, 85
Cholesky factor algorithm, 555
Cholesky factorization, 203
Cholesky-filtered log-returns, 555, 565–567
Cholesky-filtered percentage returns, 562–564
Classical rescaled adjusted range (RAR) technique, 214
Clean spread options, 646, 647
valuation of, 652–653
Close-to-maturity contracts, 272
CME, *see* Chicago Mercantile Exchange
CME WTI CSO open interest, 7
Coal Mining sector, 364
Cocoa, agricultural commodities, 485
Coffee, agricultural commodities, 484
CoinIndex (y7), business cycle variables, 20
Cointegration prices of spread, 73–76
Cointegration tests, 74–75, 207
crack spread: heating oil/WTI crude oil, 82–84
location spread: Brent/WTI crude oil, 87–88
Combined equity, performance of, 465, 466
Commission level, impact of, 294–297
Commodities, 10
agricultural, 482–485
database for commodity prices, 473, 474
energy, 479–481
hedging, 32
industrial metals, 476–479
market dependencies using wavelet squared
coherency, 390–391
performances, economic regime, 493–497
precious metals, 474–476
price, 3
Commodity calibrations, merging interest rate and,
531–534
Commodity fund, 465, 466
Commodity futures, 667
comparison of two modelling approaches for,
245–250
exchanges, 9–10
markets, 459, 462
phenomenon characteristic of, 684
portfolio choice with, 405–407
Commodity Futures Trading Commission (CFTC),
440
data, 5

- Commodity hedge funds, 461
- Commodity indexes, 391–392
- Commodity investments, 446, 459
- Commodity LIBOR Market Model
 - commodity market, 526–527
 - interest rate market model, 525–526
 - performance and applicability of, 534–541
- Commodity markets, 110, 215, 245, 434, 502, 514–515, 526–527
 - as asset class, *see* Economic regimes
 - global macroeconomic environment and, 469–470
 - statistical impact of news over different sectors of, 486, 487
 - uncorrelation of, 470
- Commodity parameters, log likelihood function for, 457–458
- Commodity portfolios, 465, 466
- Commodity prices, 96–97, 102, 106, 118, 501–503
 - commodity and equity markets, 514–515
 - conventional term-structure models of, 249
 - conventional two-factor term-structure model of, 248–249
 - cross-commodity market analysis, 512–514
 - data and empirical results, 506–512
 - dynamics, 113, 246
 - economic implications, 515–519
 - modelling, 470, 473, 474
 - models, 32
 - models and estimation, 503–506
 - short-term swings in, 10
 - stochastic dynamics of, 244
- Commodity Research Bureau, 506
- Commodity spot prices, 72
- Commodity spreads, 72
- Commodity strategies, simplified
 - combination of, 465
- Commodity Trading Advisor (CTA), 120
- Commodity-trading strategies, 460
- Common long-term trend factor models, 101
 - estimation results, 104–110
 - theoretical models, 101–104
- Comovement box, 236, 237
- Composite model, 250, 262–272, 274
- Concrete two-fuel model, 653–654
- Conditional Autoregressive Value at Risk (CAViaR) model, 237–238
- Conditional CCAPM models, 321, 322
- Conditional copula model, 408
 - in-sample results for copulas, 416–417
 - in-sample results for marginal distributions, 414–416
 - optimal portfolio results, 419–427
 - out-of-sample parameter forecasts, 417–419
- Conditional mean, 408
- Conditional parameters
 - of conditional skewed t copula, 419
 - of marginal distribution models, 418
- Conditional skewed t copula, conditional parameters of, 419
- Conditional variance of latent factor, 247
- Confidence interval, 65, 127–128
- Conservative single-step approaches, 120
- Constant conditional correlation (CCC), 419
- Constant mean reversion model
 - cumulative hedging error rates of delta hedge, 44
 - in Out of Sample WTI, 45–46
 - parameters estimation, 36–37
- Constraints, 157–159
- Consumption commodities, 244, 245
- Consumption–wealth ratio, 310
- Contango pattern, 440
- Continuous-time Gaussian factor models, 502
- Continuous-time models, advantages of, 515
- Continuous wavelet transform, 391
 - for oil futures closing prices, 351
 - for S&P500 futures closing prices, 350
- Convenience benefit, 4
- Convenience yield, 18–19
- Conventional term-structure models, 249, 250
- Conventional two-factor model, 250, 251, 260–262, 271, 272
- Conventional two-factor term-structure model, 245
 - of commodity prices, 248–249
- Conventional wisdom, 338
- Convexity correction, 530–531
- Convolution, 391
- Copper market, 440
 - average values for futures on, 441
 - futures term structure, 42
 - industrial metals, 479
 - Kalman filter state variables estimation for, 44
 - parameters for, 445
 - In Sample and Out of Sample hedging error rate, 45
 - spot price and convenience yield for, 443, 444
 - three-factor model, 39
 - two-factor model, 40
- Copula function model, 404, 409–410, 430–432
- Copulas, in-sample results for, 416–417
- Corn
 - agricultural commodities, 482, 483
 - composite model, 262, 267
 - conventional two-factor model, 260, 263
 - flexible model, 254–255
 - optimal hedge ratio for, 274, 275
- Correlation matrix, crude oil and S&P500, 341, 343
- Correlations of commodities, 469–470
- Cospectral densities, crude oil and S&P500 futures, 348–350
- ‘Cost of carry’ mechanism, 32
- Cotton, agricultural commodities, 485
- Crack-spread options, 72, 97, 121, 134
 - heating oil/WTI crude oil, 82–86

- Crack-spread option valuation, 110–111
 - commodity price dynamics, 113
 - data, 111–112
 - option valuation methodology, 112
 - results, 113–114
 - Crack-spread series, 118
 - Cramer–von Mises statistics, 558
 - CreditSpread (y_5), financial variables, 20
 - Cross-commodity market analysis, 512–514
 - Cross-contract correlation, 274
 - Cross-correlation matrix, 539, 540
 - crude oil and S&P500, 341, 344
 - Cross-correlations, 352–354
 - Cross-market price discovery research, 194, 203
 - Cross-spectral analysis, crude oil and S&P500 futures, 345–346
 - Cross wavelet phase angles, 347–348
 - Cross-wavelet power, 391
 - Cross-wavelets, 347
 - CRSP database, *see* Centre for Research in Security Prices database
 - Crude oil
 - composite model, 262, 268, 270
 - contract for, 272
 - conventional two-factor model, 260, 264
 - flexible model, 255–256
 - markets, 121–122
 - optimal hedge ratio, 274
 - portfolio variance for, 276
 - Crude oil futures
 - cospectral densities, 348–350
 - data, 341–344
 - literature review, 339–341
 - methodology, 345–348
 - wavelets, 350–354
 - CSOs, *see* Calendar spread options
 - CTA, *see* Commodity Trading Advisor
 - Cumulative distribution of VPP valuation, 583
 - Cumulative forward price distribution, 579
 - Cumulative hedging error rates
 - of delta hedge, 44, 45, 48
 - of parallel hedge, 43
- D**
- Dalian Commodity Exchange (DCE), 191
 - soybean futures, 192–194, *see also* Soybean futures
 - Dark spreads option, 646
 - Database
 - for commodity prices, 473
 - of news, 472
 - Data snooping bias, 120, 121, 124–129
 - DCC, *see* Dynamic conditional correlation
 - DCE, *see* Dalian Commodity Exchange
 - Default-free yields, 438
 - Delay parameter value, 386
 - Delta-hedge, 41, 626
 - stability of, 45–48
 - strategy derived from, 634–635
 - Demand process, 654
 - choice of parameters, 655
 - Demand variables, 21
 - Depository Institutions (DI) sector, 364, 379
 - Derivative of Gaussian wavelet, 347
 - Deseasonalization technique, 556
 - Deterministic return, ratio of, 120
 - Deviance Information Criterion (DIC), 511
 - Dickey–Fuller test, 282, 364
 - Dickey–Fuller unit root test, 60
 - Discrete jump process, 573
 - Discrete-time framework, 647
 - Discrete-time GARCH-type framework, 503
 - Discrete-time model, 601
 - Discrete-time stochastic process, 572
 - Discrete wavelet transform (DWT), 347
 - DI sector, *see* Depository Institutions sector
 - Distributed lags, dynamics of, 312
 - Diversification, 392
 - Dojima rice futures market, 459
 - Doléans–Dade exponential, 87–88, 629
 - Doléans exponential, 525
 - Dominant market, 194
 - Double-Threshold Fractionally Integrated GARCH (DT-FIGARCH) model, 359–362
 - asymmetry, tests of, 369–376
 - data and methodology, 362–369
 - framework, 361
 - model specification, tests of, 369
 - sector analysis, 376–380
 - volatility equation, 380–383
 - ‘Double whammy effect,’ 340
 - Dow Jones–UBS (DJUBS) commodity index, 6, 462
 - price patterns of, 463
 - DT-FIGARCH model, *see* Double-Threshold Fractionally Integrated GARCH model
 - d -variate copula, 407
 - d -variate Gaussian copula, 431–432
 - d -variate skewed t copula, 431
 - DWT, *see* Discrete wavelet transform
 - Dynamic conditional correlation (DCC), 419
 - model, 427
 - Dynamic investment strategy, advantage of, 434
 - Dynamic market price of risk, 450
 - Dynamic programming principle, 601–602
 - Dynamics of commodity prices
 - commodity and equity markets, 514–515
 - cross-commodity market analysis, 512–514
 - data and empirical results, 506–512
 - economic implications, 515–519
 - models and estimation, 503–506
 - Dyn model, 443

E

- Econometrics model, 58–60
- Economic regimes
 - business cycle measurement, 487–488
 - commodity markets relation to business cycle, 488–493
 - commodity performances depending on nature of, 493–497
- Economic significance of Double-Threshold FIGARCH, 379–380
- Economic variables, 10
- ECTs, *see* Error correction terms
- ECX, *see* European Climate Exchange
- EEX, *see* European Energy Exchange
- Electric-Gas Services sector, 380
- Electricity market, 571, 585–587, 624–625
 - model for spot prices, properties of, 588–593
 - option pricing, 593–600
 - price setting in, 647–648
 - pricing swing options, 600–604
- Electricity price returns, 551–553, 565, 568
 - data and preliminary analysis, 553–556
 - fitting the empirical probability densities, 558–561
 - probability distributions of, 556–558
 - robustness, 562–565
- EM, *see* Expectation Minimization algorithm
- Emission markets, 656–657
- Emission stack
 - choice of parameters, 655
 - concrete two-fuel model, 654
- Emissions Trading Scheme (ETS), 648
- EMM, *see* Equivalent martingale measure
- Empirical application
 - conditional copula model, estimation of, 417–419
 - data and preliminary analysis, 412–414
 - optimal portfolio results, 419–427
- Empirical assessment of strategies
 - data description, 440
 - optimal strategies, 446–448
 - parameters and state variables inference, 440–446
- Empirical autocorrelation function, 574
- Empirical probability densities, 558–561
- ENC-NEW tests, 314, 315, 334
- Energy derivatives, hedging strategies for, 623–626
 - literature, 626
 - non-traded assets, application to, 634–637
- Energy expenditure, 608
- Energy markets, 479–481
 - crack spread, 82–86
 - location spread, 87–90
- Engle–Granger two step test, 74, 82, 83
- EP, *see* Exponential Power
- Equal volume hedge, 81, 86
- Equity
 - and bond markets implied parameters, 443
 - market, 514–515
 - optimal proportions for, 446, 447
 - parameters, log likelihood function for, 457
 - returns, non-normality of, 502
 - time series, 438
- Equivalent martingale measure (EMM), 33
- Error correction terms (ECTs), 174
- ETS, *see* Emissions Trading Scheme
- EU, *see* European Union
- EUAs, *see* European Union Allowances
- EU ETS, *see* European Union Emission Trading Scheme
- Euler discretization, 505
- European Climate Exchange (ECX), 670
- European Energy Exchange (EEX), 572
- European gas market, 625
- European options, 72
- European Union Allowances (EUAs) returns, stylised
 - facts of, 667
 - of asset returns, 668–669
 - commodity features, 683–685
 - correlation facts, 676–678
 - data, 669–671
 - data distribution, 671–676
 - statistical properties of, 668
 - volatility-related features, 678–683
- European Union Emission Trading Scheme (EU ETS), 648, 668
- Eurostat, 684
- Ex ante market data analysis, 75
- Excess kurtosis, 54
- Excess return of optimal strategies, 448, 449
- Exchanged-traded natural gas, 160
- Exchange rates, 231, 235–236
- Exogenous stochastic factors, concrete two-fuel model, 654
- Expansionary regimes, 494
- Expansion techniques, disadvantage of, 153
- Expectation Minimization (EM) algorithm, 488, 489
- Explicit-implicit finite difference scheme, 615
- Exponential GARCH (EGARCH) model, 471, 681
- Exponential jump size distribution, 593
- Exponential Power (EP)
 - distribution laws, 552, 557
 - estimation results for, 559–561

F

- F*-adapted stochastic process, 650
- Fama–French argument, 368
- Fama–French factors, 362, 371
- Fama–French model, 369
- Fama–French three-factor model, 380
- Familywise error rate (FWER), 125, 129
- Fast Wiener–Hopf factorization (FWHF), 608
- FBSDE, *see* Forward Backward Stochastic Differential Equation

FEVD, *see* Forecasted error variance decomposition
 Feynman–Kac theorem, 614
 FIGARCH volatility equation, 380–383
 Financial assets, 667
 Financial integration theory, 470
 Financial investors, 404
 Financial market, oil price and, 304–305
 Financial variables, 20
 Finite difference scheme for pricing swing options, 613–615
 First-order serial correlation, 310
 Five-day Hui-Heubel liquidity ratio, 131, 132
 Five-day turnover rate, 131–132
 Flexible model, 252, 258, 271, 272
 of commodity futures returns, 246–248
 crude oil, 259
 gold, 260
 natural gas, 258–259
 Flexible parametric function, 248
 Föllmer–Schweizer (FS) decomposition, 627, 628
 Forecasted error variance decomposition (FEVD), 25
 structural vector auto-regression (SVAR) model, 203–205
 Vector Error Correction (VEC) model, 208–210
 Forest tree, 608
 Forward Backward Stochastic Differential Equation (FBSDE)
 for allowance price, 650–651
 numerical solution of, 663–665
 Forward price model, 576, 577
 Forwards integration method, 581, 582
 Fourier inversion, 595
 French wine index, 176
 Frequency-domain techniques, 337, 339
 Front-month contract, DCE and CBOT soybean futures, 195
 Fuel price
 changes, 657–658
 processes, 654, 655
 Futures spreads, 75–76
 spread process, 78–79
 Futures term structure of oil, 11, 12
 FWER, *see* Familywise error rate
 FWHF, *see* Fast Wiener–Hopf factorization

G

Galtchouk–Kunita–Watanabe (GKW) decomposition, 627, 628
 GARCH, *see* Generalized Autoregressive Conditional-Heteroskedasticity
 Garman and Klass volatility, 678
 Gas market, 625
 Gasoil, energy commodities, 480
 Gasoline vs. WTI crack-spread options, case of, 111
 Gaussian copula, 409–410
 Gaussian data distribution, 668, 671

Gaussian dependence, 414
 Gaussian distribution, 406
 GBM, *see* Geometric Brownian motion
 Generalized Autoregressive Conditional-Heteroskedasticity (GARCH)
 model, 59, 60, 228, 229, 233, 360, 362, 681
 process, 382
 Generalized Autoregressive Conditionally Heteroskedastic (GARCH)-based financial time series model, 280
 Generalized Autoregressive Conditionally Heteroskedastic (GARCH)-jump model, 54
 Generalized hyperbolic distribution, 552
 Generalized Method of Moments (GMM), 319–322
 Geometric Brownian motion (GBM), 10, 32, 102, 152, 244, 638, 658–659
 Geometric Ornstein Uhlenbeck (GOU), 10
 German IFO survey, 472
 German ZEW survey, 472
 Gibson–Schwartz two-factor model, 82
 implied convenience yields and difference for, 14
 Girsanov density process, 629
 Girsanov’s theorem, 525
 GKW decomposition, *see* Galtchouk–Kunita–Watanabe decomposition
 Global macroeconomic environment, 469–470
 GLOBEX electronic trading system, 215
 GMM, *see* Generalized Method of Moments
 Goldman Sachs Commodity Index (GSCI), 459, 462, 474
 price patterns of, 463
 Gold market
 average values for futures on, 441
 composite model, 262, 269
 conventional two-factor model, 260, 265
 flexible model, 257–258
 optimal hedge ratio for, 274
 parameters for, 445
 precious metals, 474–475
 spot price and convenience yield for, 443, 444
 US dollars, *see* US dollar (USD) index
 Gold–silver spread returns
 data, 215–217
 introduction, 213–215
 rescaled adjusted range analysis, 217–219
 trading implications, 219–223
 Goodness-of-fit criteria, 561
 Goodness-of-fit test, 160, 162, 565
 GOU, *see* Geometric Ornstein Uhlenbeck
 Great Depression, 306
 Greenhouse gas emissions, objectives for, 667
 Grid approach, 601–602
 GSCI, *see* Goldman Sachs Commodity Index

H

Hamilton methodology, 493
 Hanna–Quinn Criterion (HQC), 62–63

Hanson solution method, 159
 Harmonized Index of Consumer Prices (HICP), 684
 Heating oil (HO)
 crack spread, 82–86
 energy commodities, 480
 vs. WTI crack-spread options, case of, 111
 vs. WTI crude oil, case of, 118
 Heat rate, concrete two-fuel model, 653
 Heaviside theta function, 557
 Hedge fund categories, 459, 460
 Hedge ratios, 63–65
 Hedging cost, 640
 Hedging error, 519
 averages and standard deviations of, 49
 comparison of, 33
 ratio, 42
 Hedging error rates, 41, 43–47
 distribution measurement of, 48–49
 Hedging instruments, 32–33
 Hedging long-term forward contracts, 46–48
 Hedging model, skewness and kurtosis impact on, 56–58
 data and preliminary analysis, 60
 econometric model selection, 60–63
 hedged portfolio characteristics, 65–67
 modelling hedge ratios, 63–65
 Hedging option, spread option pricing and, 81–82
 Hedging performance, theoretical valuation of, 274, 276
 Hedging pressure theory, 434
 HedgingPressure (y11), trading variables, 21
 Hedging services, supply and demand for, 3
 Hedging strategies for energy derivatives, 623–626
 literature, 626
 non-traded assets, application to, 634–637
 numerical results, 637–642
 Hedging techniques, 40–43
 error rates, 43–45
 Heteroskedasticity approach, 232
 HICP, *see* Harmonized Index of Consumer Prices
 High-frequency statistical arbitrage, 280
 Historical optimal (HO) method
 impact of commission, 294–297
 prediction accuracy comparison, 290–292
 rate of return comparison, 292–293
 HO, *see* Heating oil
 HOCLSpread (y8), demand variables, 21
 HO method, *see* Historical optimal method
 HQC, *see* Hanna-Quinn Criterion
 Huber M-estimator, 674
 Hui-Heibel liquidity ratio, 131–132
 Hurst coefficient, 213, 214, 219, 222
 Hydropower assets, 581
 Hypotheses concerning functional form, 370
 Hypothesis test, 120, 128–130

I

ICE, *see* Intercontinental Exchange
 Idiosyncratic errors, 258
 IEA, *see* International Energy Agency
 IGARCH process, 382
 Ill-posed option inverse problem, 155
 Ill-posed volatility calibration problems, 155
 Implicit-explicit finite difference scheme, 665
 Implied convenience yields
 for Gibson-Schwartz (1990) two-factor model, 14
 principal components of, 15
 Implied risk-neutral probability density function (PDF), 151, 152
 comparison of methods and determination of, 160–166
 option inverse problem, 156–159
 from option prices, estimation of, 152–154
 representation of option prices and parameterization of, 156–157
 Implied volatility, 597–598
 Impulse response analysis
 structural vector auto-regression (SVAR) model, 201–203
 Vector Error Correction (VEC) model, 208–210
 Impulse-response function, 233
 Industrial metals commodities, 476–479
 Industrial Production Index (IP), 308, 310
 Inference functions for margins method, 411
 Inflation hedge, 669, 684, 685
 In Sample parameters, 48
 In-sample predictability of stock returns, 316
 In-sample predictive ability, 109
 In-sample predictive performance, 313
 In-sample results
 for copulas, 416–417
 for marginal distributions, 414–416
 Intercontinental Exchange (ICE), 60
 Interest rate market model, 525–526
 and commodity calibrations, merging, 531–534
 Interest rate parameters, log likelihood function for, 457
 Intermittency, 668, 674
 International Energy Agency (IEA), 360
 Inter-temporal arbitrage, 244
 Inter-temporal price linkage, 244–245
 ‘In-the-money’ options, 657–659
 Inventory (y9), demand variables, 21
 Inverse leverage effect, 414
 Inverse theory, 157
 in calibration of option pricing models, 152, 154–155
 Inversion model, 162
 Inversion techniques, 152
 Investment ratios, 423–425
 Investors, types of, 110

INVSP model, 160–163, 166
 IP Index, *see* Industrial Production Index
 IPSA index, 176

J

Jarque–Bera normality test, 60, 61
 Jarque–Bera statistics, 364, 671
 Jarque–Bera tests, 413
 Jarrow and Rudd's Edgeworth expansion (JR) model, 162
 JD model, *see* Jump diffusion model
 Johansen approach, 181
 Joint model, 102
 with common long-term trend for commodities, 106, 108
 without common long-term trend for commodities, 105
 JR model, *see* Jarrow and Rudd's Edgeworth expansion model
 J_T tests, 319–322
 Jump dependence parameter, 511
 Jump diffusion (JD) model, 162, 163, 166

K

Kalman filter, 11, 19, 20, 33, 43–44
 state variables vs. simultaneous equation-based state variables, 44
 technique, 116, 118
 Kelly investment strategy, 462
 Keynes–Hicks 'theory of normal backwardation,' 4, 6
 k -FWER, 125, 126
 Kolmogorov–Smirnov statistics, 558
 Kolmogorov–Smirnov tests, 413
 KPSS test, 207
 Kronecker delta density, 163–165
 Kurtosis, 54, *see also* Hedging model, skewness and kurtosis impact on
 measures of, 153, 154, 163, 166

L

Lagrange multiplier (LM) statistics, 413
 Laplace transform methods, 595
 Latent multi-factor model, 73
 Latent spot spreads, 73
 Law of one price, 73
 Lawson solution method, 159
 Lead, industrial metals, 479
 Least Squares Monte Carlo, 112
 Least-squares regression modelling, 574
 Levenberg/Marquardt algorithm, 539
 Lévy density, 608
 Lévy–Khintchine formula, 609
 Lévy models, 155
 Lévy processes
 definitions, 608–609

Regular Lévy Process of Exponential type (RLPE), 610–611
 Wiener-Hopf factorization, 611–612
 Lévy's inversion theorem, 595
 LIBOR (y_3), financial variables, 20
 LIBOR Market Model (LMM), 523–525
 calibration with time-dependent volatilities, 527–530
 commodity, 525–527
 futures/forward relation and convexity correction, 530–531
 merging interest rate and commodity calibrations, 531–534
 pricing spread option, 541–545
 (Log) likelihood functions, 456–458
 Likelihood ratio (LR) tests, 416
 statistics, 417
 Linear cointegration approach, 174, 182
 Linear model, 233
 Liquidity
 measurement, 142–143
 remark on, 131–132
 Ljung–Box statistics, 413
 Ljung–Box test, 556
 statistics, 364
 LME copper, 37
 LMM, *see* LIBOR Market Model
 Locally risk-minimizing (LRM) hedging strategy, 624, 626–627
 determination of, 630–631
 existence of, 627–629
 stochastic volatility models, 631–634
 Location spreads, 72, 121
 Brent/WTI crude oil, 87–90
 Logarithmic transformation, 562
 Log dividend–price ratio, 309–310
 Long futures-curve tactic, wealth paths for, 464, 465
 Long-horizon predictability of stock returns, 316–318
 Long-horizon regressions, 316
 Long-memory processes, 214
 'Long-only' investors, 3
 Long-run equilibrium, 174
 Long-short vs. long-only commodity funds
 alpha in commodity markets, sources of, 459–461
 relative-value commodity approach, 462–466
 temporal performance measures, 461–462
 Long-term process, 576–577
 Long-term variation process, 573
 LRM hedging strategy, *see* Locally risk-minimizing hedging strategy
 LR tests, *see* Likelihood ratio tests
 Lyapunov exponents, 283
 chaotic properties testing using, 289

M

MacDonald function, 557
 Macroeconomy, oil price and, 303–304

- Marginal distribution models
 - conditional parameters of, 418
 - in-sample results for, 414–416
 - Marginal emission rate, concrete two-fuel model, 653
 - Margins method, inference functions for, 411
 - Margrabe's classical spread option formula, 647
 - Market emission rate, allowance certificate, 649
 - Market measure
 - spot prices dynamics, 15–16
 - spread mean reversion tests, 75–76
 - spread process in, 76–78
 - Market-neutral combination of momentum, 465
 - Markov chain Monte Carlo (MCMC) estimation
 - approach, 505, 512
 - Markovian market model, 632
 - Markov model, 465
 - Markov processes, 587
 - Markov regime-switching models, 470, 488, 493–494
 - Markov-switching Vector Error Correction models, 180
 - Markowitz model, 423
 - Markowitz strategy, 419
 - Martingale Representation Theorem, 651
 - Maximum drawdown, 461
 - Maximum likelihood (ML) approach, 410–411, 416
 - Maximum likelihood estimation (MLE) optimisation
 - procedure, 55
 - MCMC estimation approach, *see* Markov chain Monte Carlo estimation approach
 - Mean pricing error (ME)
 - for crack spread, 84, 85
 - for location spread, 89
 - Mean reversion of spread, 73–76
 - Mean-reverting log spot price, 10
 - Mean-reverting (MR) process, 10, 11, 73–76, 244, 604
 - Means of Absolute values Deviating from unity (MAD), 63–65
 - Mean squared forecasting errors (MSE), 314, 334
 - Mean-variance hedging (MVH) strategy, 625
 - Mean-variance trade-off process, 629
 - Measurement equation, 28, 79, 116
 - Mediobanca Global Wine Industry Share Price Index, 174–180
 - Medium expansion regime, 496
 - Merton-Breeden hedging terms, 437
 - Metal futures markets, 280
 - Metallgesellschaft's parallel hedging, 33
 - Metropolis algorithm, 505
 - Minimal martingale measure (MMM), 628
 - Minimum-variance (MV) hedging strategy, 53
 - Mixture of two lognormals (MLN) model, 162, 163, 166
 - ML approach, *see* Maximum likelihood approach
 - MLE optimisation procedure, *see* Maximum likelihood estimation optimisation procedure
 - MLN model, *see* Mixture of two lognormals model
 - MMM, *see* Minimal martingale measure
 - Model interpretation
 - explanatory variable specification, 20–21
 - p factor impulse responses, 24–25
 - SVAR model statement, 21–22
 - x factor impulse responses, 22–23
 - y factor impulse responses, 23–24
 - Model resolution matrix, 156, 158
 - Model selection test, 252
 - Modern commodity markets, 120
 - Modern portfolio theory, 389
 - Moment generating function, 595
 - pointwise convergence of, 591
 - of truncated spike process, 590–591
 - Monte-Carlo calculation of option prices, 666
 - Monte-Carlo methods, 601, 608
 - Monte-Carlo price, 543
 - Monte-Carlo simulations, 405, 411, 581, 593, 599, 600
 - Morlet wavelet, 346, 390–391
 - Morozov discrepancy principle, 155
 - MRA, *see* Multiresolution analysis
 - MR process, *see* Mean-reverting process
 - MSE, *see* Mean squared forecasting errors
 - Multi-factor model based hedging, 33
 - Multi-period portfolio model, 461
 - Multiple hypothesis testing, 121, 124–125
 - balanced stepdown procedure, 127–129
 - single-step procedure, 125
 - stepdown procedure, 125–127
 - Multiple optimal stopping problem for swing options, 612–613
 - Multiresolution analysis (MRA), 346
 - Multivariate analysis, 413–414
 - Multivariate conditional copula with asymmetry, 407–408
 - model estimation and portfolio optimization, 410–412
 - modelling copula functions, 409–410
 - modelling univariate processes, 408–409
 - Multivariate density model, 404
 - Multivariate distribution, 410
 - Multivariate variance gamma process, 641
 - MV hedging strategy, *see* Minimum-variance hedging strategy
 - MVH strategy, *see* Mean-variance hedging strategy
- ## N
- NAP, *see* National Allocation Plan
 - National Allocation Plan (NAP), 648
 - for phases, 670
 - National Bureau of Economic Research (NBER) phase
 - of business cycle, 306
 - agricultural markets impacts depending on, 483–484
 - energy markets impacts depending on, 481
 - industrial metals impacts depending on, 478
 - precious metals impacts depending on, 476
 - Natural gas
 - composite model, 262, 266

- conventional two-factor model, 260, 261
 - energy commodities, 480
 - flexible model, 253–254
 - market, 160–166
 - optimal hedge ratio, 274
 - optimal portfolio, 272
 - portfolio variance for, 276
 - NBER, *see* National Bureau of Economic Research
 - Net hedging pressure, 5
 - Neural networks, 280
 - Newey–West estimator, 319–322
 - Newey–West long-run variance (HAC) estimator, 494
 - ‘New World Countries,’ 173–175
 - New York Mercantile Exchange (NYMEX), 97, 99, 111, 215, 251, 362–363
 - natural gas, 160, 166–167
 - New York Stock Exchange, 194
 - N*-factor mean reversion model, 32
 - Nickel, industrial metals, 479
 - NIG, *see* Normal Inverse Gaussian
 - NNLS method, *see* Non-negative least-squares method
 - Non-linear Levenberg/Marquardt algorithm, 539
 - Non-negative least-squares (NNLS) method, 151, 159
 - Non-normality of equity returns, 502
 - Non-standard financial contracts, 580
 - Non-synchronicity bias, 122
 - Non-traded assets, 626
 - application to, 634–637
 - Non-trading hour returns, DCE and CBOT soybean futures, 196–197
 - structural vector auto-regression (SVAR) model, 198–205
 - NordPool market, 587
 - Normal Inverse Gaussian (NIG)
 - distribution laws, 552
 - estimation results for, 559–561
 - Null hypothesis, 125, 219, 427
 - of DQ test, 238
 - NYMEX, *see* New York Mercantile Exchange
 - NYMEX WTI (light sweet crude oil), 37
- ## O
- Objective constraints, 157
 - Observable variables, 28
 - Observational errors, 36, 37
 - ODE, *see* Ordinary Differential Equations
 - Oil and Gas Extraction sector, 364
 - Oil-consuming sectors, 377–379
 - Oil–Gas Extraction sector, 369, 371, 379
 - Oil market
 - average values for futures on, 441
 - parameters for, 445
 - spot price and convenience yield for, 443, 444
 - Oil price
 - business cycle and stock returns, 308–309
 - cost/benefits analysis, 371
 - effects, 367
 - and financial market, 304–305
 - and macroeconomy, 303–304, 376
 - oil-consuming (substitute) sectors, 368
 - oil futures open interest term structure, 11–13
 - shocks measurement, 307–308
 - single-factor convenience yield models, 13–15
 - standard deviation (SD) of, 363
 - and stock markets, nonlinear relationship between, 340–341
 - volatility, 340
 - Oil-related sectors, 377
 - Oil-substitute sectors, 376–377
 - ‘Old World Countries,’ 173–175
 - OLS regressions, *see* Ordinary Least Squares regressions
 - Omega ratio, 423
 - One-factor model, 32
 - OPEC oil embargo of 1973, 360
 - Open interest, 11
 - vs.* time to maturity, 12
 - OpenInterest (y10), trading variables, 21
 - Open-outcry trading, 215
 - Optimal dispatch strategy modelling, 581
 - Optimal forecasting horizon for various asset classes, 496
 - Optimal hedge ratio, 63, 64, 271, 272
 - for corn, 274, 275
 - Optimal hedging strategy, implications for, 270–276
 - Optimal holding period, 219
 - Optimal portfolio proportions, 448
 - Optimal portfolio results, 419–420
 - investment ratios and portfolio performance, 423–425
 - portfolio weights, 420–422
 - robustness analysis, 425–427
 - Optimal price-dependent hourly bids, 582, 583
 - Optimal statistical arbitrage trading model, 122–124, 129
 - Optimal strategies
 - design, 435–440
 - optimal proportions for equity, bond, 446, 447
 - Sharpe ratio of, 448, 449
 - wealth of, 451
 - z* statistic for comparison Sharpe ratio of, 450
 - Option hedging, 517–519
 - Option inverse problem, 155
 - formulation and solution techniques of, 156–159
 - Option prices, 516, 593–594
 - delivery period, pricing options on forwards with, 597–600
 - for implied risk-neutral PDF, 151, 152, 154
 - Monte Carlo calculation of, 666
 - pricing options on forwards, 595–597
 - pricing path-independent options, 595
 - of probability density function (PDF), 156–157
 - semi-analytic formulae for, 586
 - Option valuation methodology, 112

Ordinary Differential Equations (ODE), 652
 Ordinary Least Squares (OLS) regressions, 307, 308
 for excess stock returns on predictor variables, 311, 313
 long-horizon predictability, 316, 317
 Ornstein–Uhlenbeck process, 76–77, 102, 120, 122, 124, 586
 Out of Sample hedges, 45–46
 Out-of-sample parameter forecasts, 417–419
 Out-of-sample predictability
 ability, 108, 110
 of stock returns, 313–315
 tests, 315, 333–335
 Out-of-sample tests, 314–315

P

Pairwise comparison, 64, 65
 Parallel hedge, 42
 hedging error rate of, 43
 Parameters and state variables inference, 440–446
 Partial integro-differential equations (PIDEs), 608, 637
 Passive indexing strategies, 462
 Paul wavelet, 347
 PCA, *see* Principal component analysis
 PCE price index, *see* Personal consumption expenditures price index
 PDF, *see* Implied risk-neutral probability density function
 Penalty function technique, 158
 Personal consumption expenditures (PCE) price index, 319
 Petroleum Refinery sector, 364, 371, 379
 PFC, *see* Price forward curve
 Philips–Perron (PP) test, 100, 181, 282, 364
 PIDEs, *see* Partial integro-differential equations
 Platinum metal, 475
 Poisson process, 504, 590
 Portfolio optimization, 410–412
 Portfolio performance, 423–426
 Portfolio selection theory, 551
 Portfolio weights, 420–422
 Positive skewness, 54
 Positivity constraints, 158–159
 Power plant value, 658
 PP test, *see* Philips–Perron test
 Precious metals commodities, 474–476, 486
 Prediction accuracy, comparison of, 290–292
 Predictive regressions, 313
 for excess stock returns, 311
 Price forward curve (PFC), 572–574
 comparison of realised and modelled price difference, 575
 and spot price model, 578
 Prices evolution, for EUA futures contracts, 670
 Pricing errors, 28
 for five model, 162

Pricing problem, allowance certificates, 649–652
 Pricing spread option, 541–545
 Pricing swing options
 in Black–Scholes model and CGMY model, 618–619
 finite difference scheme for, 613–615
 using Wiener–Hopf approach, 615–618
 Principal component analysis (PCA), 15
 Probability density function, 557
 of log-returns, 552
 Probability distribution of electricity price returns, 556–558
 Prominent commodities, 463
 Purchasing power parity, 73

Q

Quadratic utility function, 54
 Quality spreads, 72
 Quantile regressions, 229
 Quantitative trading
 in crude oil and refined products market, 120
 in oil-based markets, 119
 of spreads, 121
 Quasi-maximum likelihood method, 367

R

Radon–Nikodým derivatives, 525, 543, 545, 626, 628, 630
 Rate of return, comparison of, 292–293
 RBOB, *see* Reformulated gasoline blendstock for oxygen blending
 Reaction of commodity markets to economic news
 commodity by commodity analysis, 473–485
 database of news, 472
 measuring impact of price discovery on asset prices, 471
 rolling analysis, 485–487
 S&P 500, US 10-year and USD, 472–473
 Reality check bootstrap test, 120, 121, 129
 Realized returns
 vs. expected returns, CAPM model, 323–326
 vs. expected returns, CCAPM model, 323, 327–329
 Real per-capita consumption series, 319
 Recessionary regimes, 494
 ‘Reduced-form’ models, comparison of, 658–660
 Refined products markets, 121–122
 Refining margin, 99, 100
 Refining process, 96
 Reformulated gasoline blendstock for oxygen blending (RBOB), 99, 104
 Regional Greenhouse Gas Initiative (RGGI), 648
 Regression analysis, 573
 Regular Lévy Process of Exponential type (RLPE), 610–611
 Rescaled log-returns, 555

Rescaled percentage returns, 562–564
 Return correlation, 72, 73
 RGGI, *see* Regional Greenhouse Gas Initiative
 Rice, agricultural commodities, 485
 Risk-free asset, 671
 Risk-free rate, 308
 Risk neutral measure, 72
 spot prices dynamics, 16
 spread mean reversion tests, 75–76
 spread process in, 76
 Risk-neutral pricing of allowance certificates, 648–649
 Risk premium, 6–8
 RLPE, *see* Regular Lévy Process of Exponential type
 RMSE, *see* Root mean squared error
 Robustness analysis, 425–427
 for electricity price returns, 562–565
 Rolling analysis, commodity markets, 485–487
 Root mean squared error (RMSE), 160–162, 319–322
 for crack spread, 84, 85
 for location spread, 89

S

Safe assets, 227, 228
 Sampling frequency, 471
 Samuelson effect, 440, 473, 528
 Satellite market, 194
 SC, *see* Schwarz Criterion
 Scale-by-scale analysis, 341
 Scatter plots, 231
 Schwartz Information Criterion (SIC), 107
 Schwartz–Smith (2000) model, convenience yields
 inferred from, 13–15
 Schwarz Criterion (SC), 62–63
 SDE, *see* Stochastic differential equations
 Seasonal function consistent, with forward curve, 594
 Sector analysis, 376
 economic significance, 379–380
 of oil effects, 376–379
 Sector returns, 379–380
 data, 362
 Self-exciting Autoregressive Model, 180
 Semi-analytic formulae for option prices, 586
 Seminal portfolio theory, 388
 Sensitivity index for market, 486
 Shanghai Futures Exchange (SHFE), 280
 Shanghai Stock Exchange (SSE) index, 176
 Shapiro–Wilk normality test, 554–555
 Sharpe ratio, 123–124, 423
 of optimal strategies, 448, 449
 SHFE, *see* Shanghai Futures Exchange
 Short-horizon predictability of stock returns, 309–313
 Short term pricing errors, 13
 Short-term process, 574–575
 Short-term volatility, 575
 SIC, *see* Schwartz Information Criterion
 Silver metal, 475
 Simulation approach, 55
 Simultaneous confidence interval, 127, 128
 Simultaneous equation-based state variables, 48
 Kalman filter state variables vs., 44
 Simultaneous jump probabilities, for
 commodities, 514
 Single-factor convenience yield models, 13–15
 Single-step procedures, 121, 122, 125
 SJC copula, *see* Symmetrized Joe–Clayton copula
 Skew, 152
 Skewed *t* copula, 409, 430
 Skewness, 55, 565, *see also* Hedging model, skewness
 and kurtosis impact on
 measures of, 153, 154, 163, 166
 portfolio choice with, 405–407
 Sklar’s theorem, 408, 409
 Sortino ratio, 423
 Soybean, agricultural commodities, 485
 Soybean futures, 192
 continuous price series construction, 194–195
 futures returns calculation, 195–197
 SVAR model, 198–205
 VEC Model, 205–210
 Spark spreads, 645–646
 S&P/ASX300 index, 176
 S&P 500 commodity index, 387–389, 398–399, 506
 brief literature review, 389–390
 data, 391–394
 empirical analysis of comovements, 393, 395–398
 impact of news on, 472–473
 stock and commodities market dependencies using
 wavelet squared coherency, 390–391
 Spearman’s rank autocorrelation coefficient,
 676–677
 Speculative trading, 120, 121
 S&P500 futures
 cospectral densities, 348–350
 data, 341–344
 financial variables, 20
 literature review, 339–341
 methodology, 345–348
 wavelets, 350–354
 S&P GSCI®, *see* Standard and Poor Goldman Sachs
 Commodity Index
 Spike process, 588–589, 591
 Spot price models, 503–504
 Spot price multi-period model, 583–584
 calibration, 573–578
 overview, 572–573
 scenario generation, 578–580
 structured product valuation, 580–583
 Spot prices, 33–34, 578–580
 approximations, 589–593
 combined process, 589
 commodities, 72
 direct modelling of, 32
 dynamics of, 15–17

- process, 586
- spike process, 588–589
- Spot price simulation model, 572, 582
- Spot spread process, 73, 75, 76, 78
 - price and hedge options on, 81
- Spread levels, 75
- Spread mean reversion tests, 75–76
- Spread options, 72, 541–545, 645–646
 - choice of parameters, 656
 - literature on, 647
 - pricing and hedging, 81–82
- Spread option valuation
 - crack spread: heating oil/WTI crude oil, 86
 - location spread: Brent/WTI crude oil, 89–90
- Spread prices, numerical calculation of, 665–666
- Spread process
 - calibration, 79–81
 - cointegrated prices and mean reversion of, 73–76
 - futures pricing, 78–79
 - in market measure, 76–78
 - in risk neutral measure, 76
- Spread term structure slope, 75
- Stable distribution parameters, 558
- Stalling regime, 495, 496
- Standard and Poor Goldman Sachs Commodity Index (S&P GSCI®), 387, 389, 391–392
 - wavelet squared coherency for, 393, 395–398
- Standard asset pricing arguments, 436
- Standard delta hedging strategy, 517
- Standard errors, 233
- Standard mean-variance framework, 404
- Standard & Poor's Goldman Sachs Commodity Index, 6
- State space form, 79
 - three-factor model in, 28–29
- State variables, 40
 - correlation coefficients implied from, 444
 - expectation and covariance matrix of, 51
- Statistical arbitrage, TGARCH-WNN
 - impact of commission, 294–297
 - mean spread with TGARCH, 281–282
 - prediction accuracy comparison, 290–292
 - rate of return comparison, 292–293
 - wavelet neural networks, 283–285
 - wavelet theory and chaotic properties of time series, 282–283
- Statistical arbitrage trading strategies, 120, 122–123
- Statistics, of log-returns, 554
- Stepdown procedure, 125–127, 133
- Stepwise procedure, 121, 131
- Stochastic differential equations (SDEs), 101, 103
 - of individual state variables, 34
 - time discretization of, 665–666
- Stochastic diffusion equation, 154
- Stochastic mean reversion model
 - cumulative hedging error rates of delta hedge, 45
 - parameters estimation, 36–37
- Stochastic process, 152, 153
- Stochastic spot price multi-period model, 571–572
- Stochastic volatility framework, 502
- Stochastic volatility models, 631–634
- Stock indexes, 391–392
- Stock market index, 174
- Stock markets, 392, 551
 - dependencies using wavelet squared coherency, 390–391
 - nonlinear relationship between oil prices and, 340–341
- Stock performance, evidence on, 176–180
- Stock returns
 - and business cycle, 305–309
 - long-horizon predictability of, 316–318
 - out-of-sample predictability of, 313–315
 - short-horizon predictability of, 309–313
- Storage operators
 - in contango market, 6
 - economic function of, 4
- Strategic commodity allocation, 433–434
 - empirical assessment of strategies, 440–448
 - optimal strategies design, 435–440
- Strike price, 72
- Structural BEKK model, 229, 232, 233
- Structural disturbance vector, 197
- Structural vector auto-regression (SVAR) model, 11, 21–22, 193
 - estimation results, 200–201
 - forecasted error variance decomposition, 203–205
 - impulse response analysis, 201–203
 - model setup, 197–200
- Structured product valuation
 - valuation and risk analysis, 580–581
 - virtual power plant valuation, 582–583
- Stylised facts of EUAs returns, 667
 - analysis of, 669–671
 - of asset returns, 668–669
 - commodity features, 683–685
 - correlation facts, 676–678
 - data distribution of, 671–676
 - volatility-related features, 678–683
- Subjective constraints, 157
- Sugar, agricultural commodities, 485
- S_U -normal distribution, 55, 58–59
- SupLM tests, *see* Supremum Lagrange Multiplier tests
- Supply and demand, 4
 - for financial barrels and net hedging pressure, 5
- Supremum Lagrange Multiplier (SupLM) tests, 181
- SVAR model, *see* Structural vector auto-regression model
- SVCJ model, 504, 510, 511
- SVJ model, 509, 511–512
- Swing contracts, 600–601

Swing options, 607
 multiple optimal stopping problem for, 612–613
 numerical results, 618–619
 Symmetric dependence, 413–414
 Symmetrized Joe-Clayton (SJC) copula, 414
 Symmetry hypotheses, 370

T

Taylor effect, 669, 677
 Technical trading rules, profitability of, 121
 TermSpread (y6), business cycle variables, 20
 Testable hypotheses, 368–369
 Texas light sweet, 97
 TGARCH-WNN statistical arbitrage
 impact of commission, 294–297
 mean spread with TGARCH, 281–282
 prediction accuracy comparison, 290–292
 rate of return comparison, 292–293
 wavelet neural networks, 283–285
 wavelet theory and chaotic properties of time series, 282–283
 TGE, *see* Tokyo Grain Exchange
 Theory of normal backwardation, 10, 19
 Theory of storage, 10, 19, 22–24, 244, 245
 Third-order Taylor series expansion, 407, 420
 Three-factor FF model, 318, 320
 realized *vs.* expected returns, 324
 Three-factor Gaussian models, 32–35
 expectation and covariance matrix of state variables, 51
 number of futures units, 52
 parameters estimation, 35–40
 Three-factor model, 503
 copper, 39
 estimated latent factor evolution in, 14
 parameter estimates of, 18
 spot prices dynamics, 15–17
 in state space form, 28–29
 two-factor convenience yield, 18–19
 vs. two factor models, 19
 WTI, 38
 Three factor (log) spot price model, 17
 Three-factor stochastic mean reversion model, 48
 Three-moment optimal hedging ratios, 65
 Threshold cointegration approach, 174, 180, 183
 Threshold GARCH (TGARCH) model, 280–281
 application of, 287–289
 sample selection, 285–287
 Threshold value, 386
 Threshold vector error correction model (TVECM), 180–181
 Tikhonov approach, 155
 Time-dependent volatilities calibration, 527–530
 Time series volatility, 678
 Time-to-maturity contract, 440
 Tokyo Grain Exchange (TGE), soybean futures, 192
 Toronto Stock Exchange, 194
 TR, *see* Five-day turnover rate; Turnover rate
 Trace and Maximum-Eigenvalue tests, 181
 Trade cycle time, 122–123
 Trading hour returns, DCE and CBOT soybean futures, 196–198
 Trading rules, 214
 Trading signals, 130
 Trading strategies, profitable, 144–147
 Trading variables, 21
 Traditional portfolio model, fundamental flaws in, 459
 Traditional portfolio theory, 404
 Transaction costs
 remark on, 131
 robustness to, 135–139
 Transition equation, 28, 79, 117
 Transportation Equipment sector, 364, 379, 380
 Tree method, 601
 Truncated spike process
 distribution of, 591
 moment generating function of, 590
 TS-Energy, 580–581
t-stats, 311, 316, 317, 319, 320
 Turnover rate (TR), 131–132
 TVECM, *see* Threshold vector error correction model
 Two-factor convenience yield, 18–19
 Two-factor Gaussian model, 243–245, 248
 composite model, 262–270
 conventional two-factor model, 260–262
 conventional two-factor term-structure model of commodity prices, 248–249
 data, 251
 flexible model of commodity futures returns, 246–248, 252–260
 model comparison, 249–250
 optimal hedging strategy, 270–276
 Two-factor mean reversion model, 32
 Two-factor model, 104, 105
 copper, 40
 delta hedge, 41
 parameter estimates for, 85, 89
 parameter estimates of, 18
 three-factor model statement *vs.*, 19
 WTI, 39
 Two-moment utility function, 54
 Two price method of simulating spreads, 92–93
 Two-stage log-likelihood function, 432

U

Ulcer index, 461, 463
 Ulcer performance index (UPI), 464
 Unconditional asset pricing models, 320
 Unconditional multivariate Gaussian model, 419
 Uncorrelated model, 113–114
 Uncorrelation of commodity markets, 470

Unit roots, 74–75
 crack spread: heating oil/WTI crude oil, 82–84
 location spread: Brent/WTI crude oil, 87–88
 Unit root tests, 100, 207
 Univariate analysis, 413, 506–512
 Univariate processes model, 408–409
 UPI, *see* Ulcer performance index
 US dollar (USD) index
 dataset, 230–232
 financial variables, 20
 impact of news on, 472–473
 overview of, 228–229
 structural evidence on volatility spillovers, 232–236
 U.S. Strategic Petroleum Reserve (SPR), 361

V

Valuation and risk analysis application, 580–581
 Value at Risk (VaR) method, 674, 678
 Value-weighted CRSP index, 308, 312
 Variance–covariance matrix, 232
 Variance gamma function, 641
 Variance Gamma Processes (VGP), 610
 Variational inequalities, treatment of, 614–615
 VaR method, *see* Value at Risk method
 VARs, *see* Vector autoregressions
 Vech-GARCH model, 233
 VEC model, *see* Vector Error Correction model
 Vector autoregressions (VARs), 340
 model, 232
 process, 75
 technique, 360
 Vector Error Correction (VEC) model, 193, 205
 estimation results, 206–208
 impulse response and FEVD, 208–210
 model setup, 206
 VGP, *see* Variance Gamma Processes
 VIRF, *see* Volatility Impulse Response Functions
 Virtual power plant (VPP)
 simplified setup of, 581
 valuation of, 582–583
 VIX (y4), financial variables, 20
 Volatility, 502
 in crude oil and refined products markets, 121
 estimation of surface, 579–580
 term structure of, 577–578
 Volatility clustering, 669, 678–680
 Volatility correlations, 513
 Volatility Impulse Response Functions (VIRF), 233
 Volatility integrals, 547–548
 Volatility-related features, 678–683
 Volatility shocks concept, 229
 Volatility smile, 152
 Volatility spillovers, structural evidence on, 232–236
 Vol-of-vol parameter, 508–510
 VPP, *see* Virtual power plant
 Vuong's test, 493, 494

W

Wald test
 bootstrapped p -values for, 316, 317
 CAPM and CCAPM models, 320–322
 short-horizon predictability of stock returns, 311–313
 Wavelet coherency methodology, 389
 Wavelet-filtered log-returns, 565–567
 Wavelet function, 390
 Wavelet neural networks (WNNs), 280, 281, 283–285
 architecture network of, 290
 Wavelet power spectrum
 oil futures, 351, 354
 S&P500 future, 352, 354
 Wavelets, 346–347
 application in finance, 341
 crude oil and S&P500 futures, 350–354
 Wavelet squared coherency, 395
 scale for, 393
 stock and commodities market dependencies using, 390–391
 Wavelet theory of time series, 282–283
 Wavelet transforms, 390, 391
 Wealth of optimal strategies, 448, 451
 West Texas Intermediate (WTI), 480
 futures term structure, 41
 three-factor model, 38
 two-factor model, 39
 West Texas Intermediate (WTI) crude oil, 97, 99, 524, 534–535, 537
 crack spread, 82–86
 location spread, 87–90
 oil futures prices and open interest for, 11
 West Texas Intermediate (WTI) index performance, 5, 6
 Wheat, agricultural commodities, 483, 484
 Wiener–Hopf approach, 608, 618
 pricing swing options using, 615–618
 Wiener–Hopf factorization, 611–612
 Wiener process, 122
 Wine market
 econometric methodology, 180–181
 empirical results, 181–185
 Mediobanca Global Wine Industry Share Price Index, 176–180
 theoretical framework, 175–176
 WNNs, *see* Wavelet neural networks
 WTI, *see* West Texas Intermediate

Z

ZA-test procedure, 363
 Zero-coupon bonds, 47
 volatilities of, 454
 Zinc calendar spread arbitrage, 285–287
 Zinc, industrial metals, 479
 z statistic for comparison Sharpe ratio, 448, 450